Bayesian Analysis Users Guide
Release 4.00, Manual Version 1

G. Larry Bretthorst
Biomedical MR Laboratory
Washington University School Of Medicine,
Campus Box 8227
Room 2313, East Bldg.,
4525 Scott Ave.
St. Louis MO 63110
http://bayes.wustl.edu
Email: larry@bayes.wustl.edu

August 21, 2013

Contents

Manual Status **14**

1 An Overview Of The Bayesian Analysis Software **17**
1.1 The Server Software . . . 17
1.2 The Client Interface . . . 20
1.2.1 The Global Pull Down Menus . . . 22
1.2.2 The Package Interface . . . 22
1.2.3 The Viewers . . . 25

2 Installing the Software **27**

3 the Client Interface **29**
3.1 The Global Pull Down Menus . . . 31
3.1.1 the Files menu . . . 31
3.1.2 the Packages menu . . . 36
3.1.3 the WorkDir menu . . . 41
3.1.4 the Settings menu . . . 42
3.1.5 the Utilities menu . . . 46
3.1.6 the Help menu . . . 47
3.2 The Submit Job To Server area . . . 47
3.3 The Server area . . . 48
3.4 Interface Viewers . . . 49
3.4.1 the Ascii Data Viewer . . . 49
3.4.2 the fid Data Viewer . . . 51
3.4.3 Image Viewer . . . 56
3.4.3.1 the Image List area . . . 56
3.4.3.2 the Set Image area . . . 58
3.4.3.3 the Image Viewing area . . . 58
3.4.3.4 the Grayscale area on the bottom . . . 60
3.4.3.5 the Pixel Info area . . . 60
3.4.3.6 the Image Statistics area . . . 60
3.4.4 Prior Viewer . . . 62
3.4.5 fid Model Viewer . . . 65
3.4.5.1 The fid Model Format . . . 65

3.4.5.2 The Fid Model Reports . . . 67
3.4.6 Plot Results Viewer . . . 68
3.4.6.1 the Data, Model and Residuals Plots . . . 70
3.4.6.2 the Posterior Probabilities Plots . . . 71
3.4.7 the Posterior Probability Vs Parameter Samples plot . . . 72
3.4.7.1 the Expected Log Likelihood Plot . . . 75
3.4.7.2 the Scatter Plots . . . 75
3.4.7.3 the Log Probability Plot . . . 78
3.4.8 Text Results Viewer . . . 80
3.4.9 Files Viewer . . . 86
3.4.10 Fortran/C Code Viewer . . . 86
3.4.10.1 Fortran/C Model Viewer Popup Editor . . . 88

4 An Introduction to Bayesian Probability Theory 91
4.1 The Rules of Probability Theory . . . 91
4.2 Assigning Probabilities . . . 94
4.3 Example: Parameter Estimation . . . 101
4.3.1 Define The Problem . . . 102
4.3.1.1 The Discrete Fourier Transform . . . 102
4.3.1.2 Aliases . . . 105
4.3.2 State The Model—Single-Frequency Estimation . . . 106
4.3.3 Apply Probability Theory . . . 107
4.3.4 Assign The Probabilities . . . 110
4.3.5 Evaluate The Sums and Integrals . . . 112
4.3.6 How Probability Generalizes The Discrete Fourier Transform . . . 115
4.3.7 Aliasing . . . 118
4.3.8 Parameter Estimates . . . 124
4.4 Summary and Conclusions . . . 127

5 Given Exponential Model 129
5.1 The Bayesian Calculation . . . 131
5.2 Outputs From The Given Exponential Package . . . 133

6 Unknown Number of Exponentials 135
6.1 The Bayesian Calculations . . . 137
6.2 Outputs From The Unknown Number of Exponentials Package . . . 140

7 Inversion Recovery 143
7.1 The Bayesian Calculation . . . 145
7.2 Outputs From The Inversion Recovery Package . . . 146

8 Bayes Analyze 147
8.1 Bayes Model . . . 151
8.2 The Bayes Analyze Model Equation . . . 153
8.3 The Bayesian Calculations . . . 159
8.4 Levenberg-Marquardt And Newton-Raphson . . . 163

8.5 Outputs From The Bayes Analyze Package . . . 167
8.5.1 The "bayes.params.nnnn" and "bayes.model.nnnn" Files . . . 169
8.5.1.1 The Bayes Analyze File Header . . . 169
8.5.1.2 The Global Parameters . . . 174
8.5.1.3 The Model Components . . . 175
8.5.2 The "bayes.output.nnnn" File . . . 177
8.5.3 The "bayes.probabilities.nnnn" File . . . 181
8.5.4 The "bayes.log.nnnn" File . . . 184
8.5.5 The "bayes.status.nnnn" and "bayes.accepted.nnnn" Files . . . 187
8.5.5.1 The "bayes.model.nnnn" File . . . 188
8.5.6 The "bayes.summary1.nnnn" File . . . 189
8.5.7 The "bayes.summary2.nnnn" File . . . 190
8.5.8 The "bayes.summary3.nnnn" File . . . 191
8.6 Bayes Analyze Error Messages . . . 192

9 Big Peak/Little Peak **197**
9.1 The Bayesian Calculation . . . 199
9.2 Outputs From The Big Peak/Little Peak Package . . . 206

10 Metabolic Analysis **209**
10.1 The Metabolic Model . . . 213
10.2 The Bayesian Calculation . . . 215
10.3 The Metabolite Models . . . 218
10.3.1 The IPGD_D2O Metabolite . . . 218
10.3.2 The Glutamate.2.0 Metabolite . . . 222
10.3.3 The Glutamate.3.0 Metabolite . . . 225
10.4 The Example Metabolite . . . 226
10.5 Outputs From The Bayes Metabolite Package . . . 228

11 Find Resonances **229**
11.1 The Bayesian Calculations . . . 231
11.2 Outputs From The Bayes Find Resonances Package . . . 236

12 Diffusion Tensor Analysis **237**
12.1 The Bayesian Calculation . . . 239
12.2 Using The Package . . . 244

13 Big Magnetization Transfer **249**
13.1 The Bayesian Calculation . . . 249
13.2 Outputs From The Big Magnetization Transfer Package . . . 252

14 Magnetization Transfer **255**
14.1 The Bayesian Calculation . . . 257
14.2 Using The Package . . . 261

15 Magnetization Transfer Kinetics **267**
15.1 The Bayesian Calculation . . . 269
15.2 Using The Package . . . 273

16 Given Polynomial Order **277**
16.1 The Bayesian Calculation . . . 279
16.1.1 Gram-Schmidt . . . 279
16.1.2 The Bayesian Calculation . . . 280
16.2 Outputs From the Given Polynomial Order Package . . . 282

17 Unknown Polynomial Order **285**
17.1 Bayesian Calculations . . . 287
17.1.1 Assigning Priors . . . 288
17.1.2 Assigning The Joint Posterior Probability . . . 289
17.2 Outputs From the Unknown Polynomial Order Package . . . 291

18 Errors In Variables **295**
18.1 The Bayesian Calculation . . . 297
18.2 Outputs From The Errors In Variables Package . . . 300

19 Behrens-Fisher **303**
19.1 Bayesian Calculation . . . 303
19.1.1 The Four Model Selection Probabilities . . . 306
19.1.1.1 The Means And Variances Are The Same . . . 307
19.1.1.2 The Mean Are The Same And The Variances Differ . . . 309
19.1.1.3 The Means Differ And The Variances Are The Same . . . 310
19.1.1.4 The Means And Variances Differ . . . 311
19.1.2 The Derived Probabilities . . . 312
19.1.3 Parameter Estimation . . . 313
19.2 Outputs From Behrens-Fisher Package . . . 314

20 Enter Ascii Model **321**
20.1 The Bayesian Calculation . . . 323
20.1.1 The Bayesian Calculations Using Eq. (20.1) . . . 323
20.1.2 The Bayesian Calculations Using Eq. (20.2) . . . 324
20.2 Outputs Form The Enter Ascii Model Package . . . 327

21 Test Your Own ASCII Model **329**

22 Ascii Model Selection **331**

23 Phasing An Image **333**
23.1 The Bayesian Calculation . . . 334
23.2 Using The Package . . . 340

24 Phasing An Image Using Non-Linear Phases **343**
24.1 The Model Equation . . . 343
24.2 The Bayesian Calculations . . . 345
24.3 The VnmrJ and Vnmr Interfaces . . . 347

28 Analyze Image Pixel **361**
28.1 Modification History . . . 363

29 Image Pixel Model Selection **365**

A Ascii Data File Formats **367**
A.1 Ascii Input Data Files . . . 367
A.2 Ascii Image File Formats . . . 368
A.3 The Abscissa File Format . . . 369

B Markov chain Monte Carlo With Simulated Annealing **375**
B.1 Metropolis-Hastings Algorithm . . . 376
B.2 Multiple Simulations . . . 377
B.3 Simulated Annealing . . . 378
B.4 The Annealing Schedule . . . 378
B.5 Killing Simulations . . . 379
B.6 the Proposal . . . 380

C Thermodynamic Integration **381**

D McMC Values Report **385**

E Writing Fortran/C Models **391**
E.1 Model Subroutines, No Marginalization . . . 391
E.2 The Parameter File . . . 394
E.3 The Subroutine Interface . . . 396
E.4 The Subroutine Declarations . . . 398
E.5 The Subroutine Body . . . 399
E.6 Model Subroutines With Marginalization . . . 400

F the Bayes Directory Organization **405**

G 4dfp Overview **407**

H Outlier Detection **411**

Bibliography **415**

List of Figures

1.1 The Start Up Window . . . 21
1.2 Example Package Interface . . . 23

3.1 The Start Up Window . . . 30
3.2 The Files Menu . . . 31
3.3 The Load Image Selection Menu . . . 33
3.4 The Packages Menu . . . 37
3.5 The Working Directory Pull Down Menu . . . 42
3.6 The Working Directory Po pup . . . 43
3.7 The Settings Pull Down Menu . . . 44
3.8 The McMC Parameters Po pup . . . 44
3.9 The Edit Server Popup . . . 45
3.10 The Submit Job Widget Group . . . 48
3.11 The Server Widget Group . . . 49
3.12 the Ascii Data viewer . . . 50
3.13 the fid Data viewer . . . 52
3.14 The Fid Data Viewer Display Type . . . 53
3.15 The Fid Data Viewer the Options Menu . . . 54
3.16 The Image Viewer . . . 57
3.17 The Image Viewer Right Mouse Menu . . . 58
3.18 The Prior Viewer . . . 63
3.19 The Fid Model Viewer . . . 66
3.20 The Data Model and Residuals . . . 69
3.21 The Plot Information popup . . . 70
3.22 The Posterior Probabilities . . . 71
3.23 The Posterior Probabilities Vs Parameter Value . . . 73
3.24 The Posterior Probabilities Vs Parameter Value a Skewed Example . . . 74
3.25 The Expected Log Likelihood . . . 76
3.26 The Scatter Plots . . . 77
3.27 The Log Probability Plot . . . 79
3.28 The Text Results Viewer . . . 81
3.29 The Bayes Condensed File . . . 84
3.30 Fortran/C Model Viewer . . . 87
3.31 Fortran/C Model Viewer . . . 88

4.1 Frequency Estimation Using The DFT . . . 104
4.2 Aliases . . . 105
4.3 Nonuniformly Nonsimultaneously Sampled Sinusoid . . . 119
4.4 Alias Spacing . . . 120
4.5 Which Is The Critical Time . . . 122
4.6 Example, Frequency Estimation . . . 123
4.7 Estimating The Sinusoids Parameters . . . 125

5.1 the Exponential interface . . . 130

6.1 the Unknown Exponential interface . . . 136
6.2 The Distribution of Models . . . 141
6.3 Exponential Probability for the Model . . . 142

7.1 the Inversion Recovery interface . . . 144

8.1 Bayes Analyze Interface . . . 148
8.2 Bayes Analyze Fid Model Viewer . . . 152
8.3 The Bayes Analyze File Header . . . 170
8.4 The bayes.noise File . . . 172
8.5 Bayes Analyze Global Parameters . . . 175
8.6 Bayes Analyze Model File . . . 176
8.7 Bayes Analyze Initial Model . . . 178
8.8 Base 10 Logarithm Of The Odds . . . 178
8.9 The bayes.output.nnnn Report . . . 179
8.10 Bayes Analyze Uncorrelated Output . . . 180
8.11 The bayes.probabilities.nnnn File . . . 182
8.12 The bayes.log.nnnn File . . . 185
8.13 The bayes.status.nnnn File . . . 187
8.14 The bayes.model.nnnn File . . . 188
8.15 The bayes.model.nnnn File Uncorrelated Resonances . . . 189
8.16 Bayes Analyze Summary Header . . . 189
8.17 The Summary2 Report . . . 190
8.18 The Summary2 Report . . . 191

9.1 The Big Peak/Little Peak Interface . . . 198
9.2 The Time Dependent Parameters . . . 208

10.1 The Bayes Metabolite Interface . . . 210
10.2 Bayes Metabolite Viewer . . . 212
10.3 Bayes Metabolite Probabilities List . . . 217
10.4 The IPGD_D20 Metabolite . . . 219
10.5 Bayes Metabolite IPGD_D20 Spectrum . . . 220
10.6 Bayes Metabolite, The Fraction of Glucose . . . 221
10.7 Glutamate Example Spectrum . . . 223
10.8 Estimating The F_{c0}, y and F_{a0} Parameters . . . 226
10.9 Bayes Metabolite, The Ethyl Ether Example . . . 227

11.1 the Find Resonances interface 230

12.1 Diffusion Tensor Interface 238
12.2 Diffusion Tensor Parameter Estimates 246
12.3 Diffusion Tensor Posterior Probability For The Model 246

13.1 The Big Magnetization Package Interface 250
13.2 Big Magnetization Transfer Example Fid 252
13.3 Big Magnetization Transfer Expansion 253
13.4 Big Magnetization Transfer Peak Pick 254

14.1 Magnetization Transfer Interface 256
14.2 Magnetization Transfer Peak Pick 262
14.3 Magnetization Transfer Example Data 263
14.4 Magnetization Transfer Example Spectrum 264

15.1 Magnetization Transfer Kinetics Interface 268
15.2 Magnetization Transfer Kinetics Arrhernius Plot 274
15.3 Magnetization Transfer Kinetics Water Viscosity Table 275

16.1 Given Polynomial Order Package Interface 278
16.2 Given Polynomial Order Scatter Plot 284

17.1 Unknown Polynomial Order Interface 286
17.2 The Distribution of Models 290
17.3 Unknown Polynomial Order Package Posterior Probability 292

18.1 Errors In Variables Interface 296
18.2 Errors In Variables McMC Values File 302

19.1 the Behrens-Fisher interface 304
19.2 Behrens-Fisher Hypotheses Tested 305
19.3 Behrens-Fisher Console Log 315
19.4 Behrens-Fisher Status Listing 316
19.5 Behrens-Fisher McMC Values File, The Preamble 317
19.6 Behrens-Fisher McMC Values File, The Middle 318
19.7 Behrens-Fisher McMC Values File, The End 319

20.1 Enter Ascii Model Interface 322

21.1 Test Your Own Ascii Model Interface 330

22.1 Ascii Model Selection Interface 332

23.1 Absorption Model Images 334
23.2 Bayes Phase Interface 335
23.3 Bayes Phase Listing 341

24.1 Nonlinear Phasing Example 344
24.2 Nonlinear Phasing Interface 348

28.1 Image Pixels Example 362

A.1 Ascii Data File Format 368

D.1 The McMC Values Report Header 386
D.2 McMC Values Report, The Middle 387
D.3 The McMC Values Report, The End 388

E.1 Writing Models A Fortran Example 392
E.2 Writing Models A C Example 393
E.3 Writing Models, The Parameter File 395
E.4 Writing Models Fortran Declarations 399
E.5 Writing Models Fortran Example 402
E.6 Writing Models The Parameter File 403

G.1 The FD File Header 409

H.1 the Posterior Probability for the Number of Outliers 412
H.2 The Data, Model and Residual Plot With Outliers 414

List of Tables

8.1 Multiplet Relative Amplitudes 157
8.2 Bayes Analyze Models 173
8.3 Bayes Analyze Short Descriptions 186

Chapter 3

the Client Interface

The interface to the Bayesian Analysis software is a Java interface that runs on any machine having Java 6 or higher. Assuming the Bayesian Analysis software has been installed on a server at your site, for arguments sake lets call this machine "your.server.net," then the client interface can be displayed by issuing:

```
javaws http://your.server.net:8080/Bayes/launch.jnlp
```

where "javaws" is the Java web-start utility and comes with most Java installations, "your.server.net" should be replaced by the server name or IP address, and "8080" should be replaced by the port number used during installation, see Chapter 2 for a description of how to install the software.

When the interface starts it will displays the default start up page shown in Fig. 3.1. This figure is a repeat of the figure shown in Chapter 3, Fig. 1.1 and is repeated here for convenience. The purpose of the start up page is to allow an analysis to be restarted. When the interface exits or changes working directories, the interface saves the current settings in a special Java properties file. When the interface start, it consults this file to determines what the last WorkDir was and how to restart that analysis. If an analysis was saved, the interface displays the messages shown in Fig. 3.1, the lines starting "To restore analysis". This line contains the name of the package that was being processed, in this case the package name was "AnalyzeImagePixels" and the analysis was saved in a WorkDir named "Given". If the "Restore Analysis" button is activated then the "Given/AnalyzeImagePixels" analysis will be restored to its previous status. When the interface finishes restoring the analysis, it will function exactly like WorkDir or interface was never exited.

If an analysis is not to be restored, then changing the package will delete the contents of the current WorkDir and configure the WorkDir for the new package. If another analysis is running in a different WorkDir, then changing the current working directory using the "WorkDir" menu will cause the interface to switch to the new WorkDir and the previous analysis will be restored.

Finally, if a completely new WorkDir is needed, then selecting "WorkDir/Edit" will bring up a popup that can create a new WorkDir. After the WorkDir is create, the first thing that must be done is to select a package.

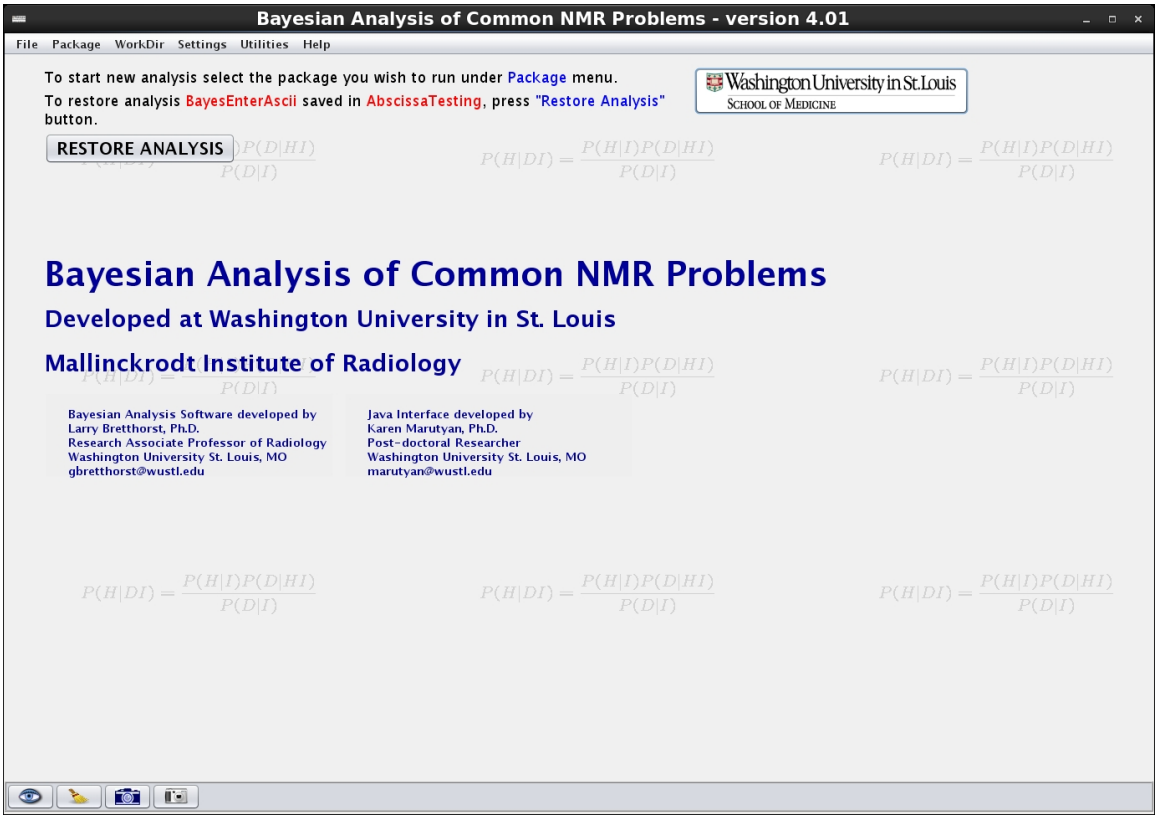
Bayesian Analysis of Common NMR Problems - version 4.01
File Package WorkDir Settings Utilities Help
To start new analysis select the package you wish to run under Package menu.
To restore analysis BayesEnterAscii saved in AbscissaTesting, press "Restore Analysis" button.
Washington University in St.Louis
School of Medicine
RESTORE ANALYSIS
Bayesian Analysis of Common NMR Problems
Developed at Washington University in St. Louis
Mallinckrodt Institute of Radiology
Bayesian Analysis Software developed by
Larry Bretthorst, Ph.D.
Research Associate Professor of Radiology
Washington University St. Louis, MO
gbretthorst@wustl.edu
Java Interface developed by
Karen Marutyan, Ph.D.
Post-doctoral Researcher
Washington University St. Louis, MO
marutyan@wustl.edu

Figure 3.1: The Bayesian Analysis Start up Page allows you to select what functions you wish to perform. For example you might restore an old analysis, change a setting, run one of the utility programs or select a new WorkDir or a new Bayesian Analysis package.

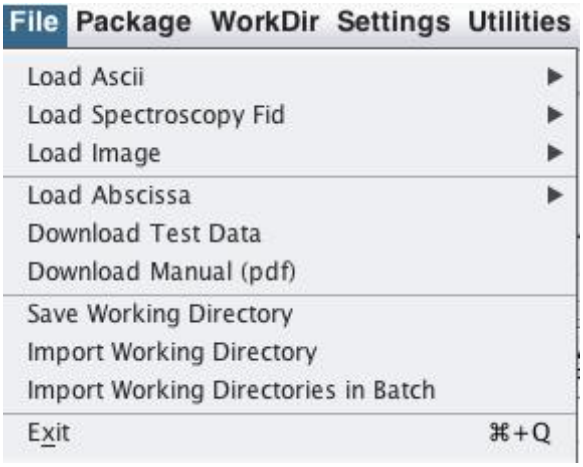

The Files menu can be used to:

Load ASCII files from multiple sources.

Load Spectroscopic fid from Varian, Siemens and ASCII files.

Load Image from ASCII, fdf, dicom, 4dfp and ima files.

Load Abscissa for image processing.

Download Test Data to the current Bayes Home directory.

Download Manual (pdf) download this manual to the current Bayes Home directory.

Save Working Directory saves the current WorkDir to a location of your choosing.

Import Working Directory reloads a saved WorkDir to a name of your choice.

Import Working Directories in Batch reloads a groups of working directories.

Figure 3.2: When the Files menu is selected, this pull down menu is displayed. Use this menu to load data, download updates to this manual and save and import working directories.

3.1 The Global Pull Down Menus

The global pull down menus along the top of the start up page are always present on all package interfaces, not just the start up page. They allow the user to load files, select packages, configure servers, change working directories, set options, etc. Each pull down menu has multiple functions and the following subsections explain these menus in detail and how to go about using them. We will take the menu from left to right starting with perhaps the most complicated menu, the Files menu.

3.1.1 the Files menu

The Files menu is a general purpose menu that handle most functions concerning loading images, ASCII data and other types of files into the Bayesian Analysis software. Figure 3.2 shows what this menu looks like when activated. Here is what each selection menu on the Files menu does:

Load ASCII selection menu loads ASCII data from either an ASCII file or a Bayes Analyze file. In either case the data is copied/reformatted and saved in the BayesOtherAnalysis directory of the current WorkDir. When multiple ASCII files are loaded the currently selected data file is plotted. The file format for an ASCII file is package specific, i.e., the number of data and abscissa columns required varies with the package. The file format for each package is addressed in the Chapter describing the package. However, the general file format used by Bayesian ASCII software is described in Section A. After the data is loaded into the WorkDir, the ASCII Data Viewer is then activated and the ASCII data is plotted in the viewer.

The Load ASCII menu has two submenus, one to load a plain ASCII file and one that extracts amplitudes from a Bayes Analyze file. Here is a description of these selection menus:

File brings up a navigation popup that navigates to the appropriate ASCII file and then load that file. Note that file is parsed by the interface to determine if it has the correct number of data columns and that it contains only ASCII data. The loaded ASCII file is copied into the current WorkDir and the file is assigned a name using a sequential number. These numbers are unique so multiple ASCII files having the same name can be loaded without conflict. The loaded file is then displayed in the ASCII Data Viewer.

Bayes Analyze File is a selection menu that will load the amplitudes of a resonance as a function of some arrayed parameter. These amplitudes are read from the previous run of the Bayes Analyze package. A prompt for the resonances number who's amplitudes are to be loaded will appear. The Bayes Analyze analysis must use a joint analysis, i.e., all fid's are analyzed jointly looking for common frequencies with fid dependent amplitudes. The amplitudes are combined with the arrayed variable from the procpar in the "fid" subdirectory in the current WorkDir.

Load Spectroscopic fid loads time domain spectroscopic fid data from several sources: Varian fid data, Siemens rda, Siemens Raw data and ASCII Text fid data. When any of these selection menus are activated, they brings up a popup that allows navigation to the appropriate file and then load it. When data are loaded, the data are copied to the "fid" Subdirectory of the current WorkDir and a Varian fid, procpar and text file are written into this Subdirectory. If the loaded fid is a Varian fid, the procpar is copied from the source directory, otherwise a procpar is generated and modified to reflect the number of data values, acquisition time and sweep width of the current data. The data are then Fourier transformed, and the real part of the discrete Fourier transform is displayed using the phasing parameters in the procpar.

Load Image menu handles the task of loading various kinds of image data into the Bayesian Analysis package. The menu will loads both k-space and image data. The first two selection menus items will read k-space data and then convert the k-space data into images. The remaining menu items all read various types of images and convert them into our internal format. Images are stored in the "Images" Subdirectory of the current WorkDir. No attempt is made to change the name of an image so loading images with the same name will result in replacing the old image. We will briefly mention the various types of data that can be loaded and give a brief description of the file formats.

Varian k-space fid will bring up a popup that allows navigation loading of standard Varian k-space image fid. The selected file must be suffixed with ".fid" and a Varian binary file is expected. The binary is copied into the "image.fid" Subdirectory in the current WorkDir. It is then Fourier transformed, phased and three ".img" file are written into the "images" Subdirectory. These three files are named "LoadedImage-Abs.4dfp.img," "LoadedImage-Real.4dfp.img," and "LoadedImage-Imag.4dfp.img" and contain the absolute value, imaginary and real images.

Text k-space fid will bring up a popup that allows navigating to and load a text k-space image fid. The selected file must be suffixed with ".fid" and a Varian binary file is expected. The binary is copied into the "image.fid" Subdirectory in the current WorkDir.

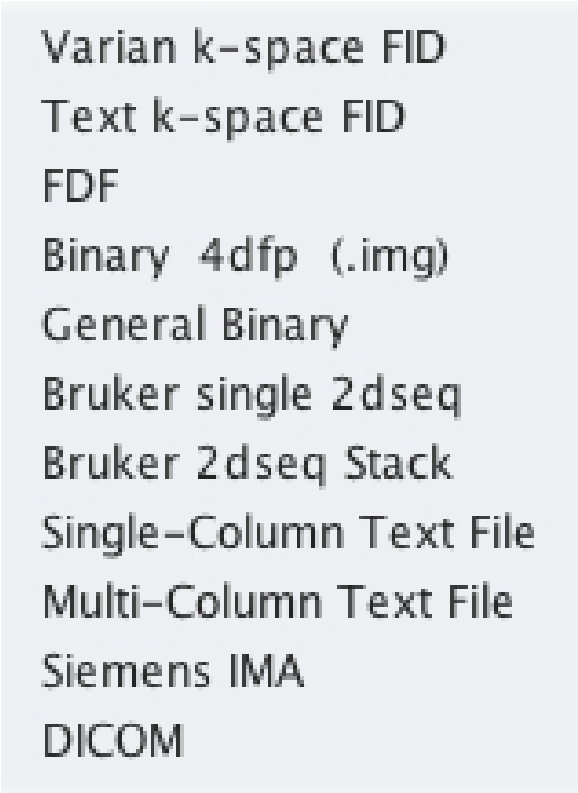

Varian k-space fid loads a Varian k-space fid and converts it to a 4dfp stack.

Text k-space fid loads an ASCII Text k-space fid and converts it to a 4dfp stack.

FDF loads and FDF stack of images and reorders and converts it to a 4dfp stack.

Binary 4dfp (.img) loads 4dfp stack of images.

General Binary brings up a popup that will attempt to load a general binary file.

Bruker single 2dseq brings up a popup loads a single Bruker 2dseq file.

Bruker 2dseq Stack brings up a popup loads a Bruker stacked 2dseq file.

Single-Column Text File brings up a popup Loads single-column text images.

Multi-Column Text File brings up a popup Loads multi-column text images.

Siemens IMA brings up a popup Loads Siemens IMA images.

DICOM brings up a popup Loads DICOM images.

Figure 3.3: When the Files/Load Image selection menu is selected this pull down menu is displayed. It is used to select the type of image data to be loaded. After selecting a data type, a popup will be displayed that will allow navigation to the appropriate loading of the data. In most cases the popup will have a number of configuration parameters that have to enter.

It is then Fourier transformed, phased and a three ".img" file are written into the "images" Subdirectory. The image files images are names as discussed in the previous item.

FDF will bring up a popup that allows navigating and loading of Varian FDF images. One or more of the FDF files are loaded by the open button. This will copy and reformat the FDF images into a single 4dfp file located in the "images" Subdirectory of the current WorkDir. The input images are ordered so that the displayed images are in the proper slice and array order.

Binary 4dfp (.img) will bring up a popup that allows navigating and loading of 4dfp file. The 4dfp file is copied into the images Subdirectory. The 4dfp file type is the internal standard in which all images are stored in the Bayesian Analysis software, so no additional processing is needed. See Appendix G for a description of 4dfp files.

General Binary

Bruker single 2dseq

Bruker 2dseq Stack

Single-Column Text File will bring up a popup that allows one to navigate to a single-column text file and load it into the images Subdirectory. The single-column text is read by the interface and then parsed into images with the assistance of a popup window. The popup will display the total number lines in the text file and the row, column, slice and array dimensions must be set so that the total pixels is equal to the total lines. Until these dimensions are set correctly the interface will not load the image. Multiple images can be stacked in the file. If so, they can be stacked either in slice or array order and specify which order is used when the data are loaded. The default outer loop, the most slowly varying loop, is the array dimension with slice as the inner loop.

Multi-Column Text File will bring up a popup that allows navigating and loading of a Multi-Column text file. Multi-Column text files (images) are read by the interface as series of lines each containing multiple pixels. Each line in the file corresponds to one horizontal line in the displayed image. For MRI data each line corresponds to the phase encode direction. The number of phase encode pixels depends on the number of phase encodes and the zero padd level of the Fourier transform. So if the image has a total of 64 phase encode pixels (including zero padding), then each line in the text file must have 64 phase encode pixels. Additionally, if there are 96 pixels in the vertical (readout) direction, then there must be 96 total lines for each image in the text file. Like single-column text files, multiple images can be stacked in one file and the images can be ordered by either by slice or array element. Also, like single-column text files, when an image is selected, a popup will be displayed and the number of slices, array dimension, and image sizes must be specified.

Siemens IMA will bring up a popup that allows navigating to loading of of Siemens IMA images. One or more of the IMA files can be selected and opened. This will copy and reformat the IMA images into a single 4dfp file located in the "images" Subdirectory of the current WorkDir. The order of the images is alphabetical, if the images must be ordered in some special way they must be renamed appropriately.

DICOM will bring up a popup that allows navigating to and loading of DICOM images. This will copy and reformat the DICOM images into a single 4dfp file located in the

"images" Subdirectory of the current WorkDir. The interface attempts to order the images internally so that the displayed images are in the proper slice and array order.

Load Abscissa is submenus on the Files menu. When selected, it brings up a popup that allows the user to navigate to an to a file and then load it as an Abscissa file. Abscissa files are stored in the "images" Subdirectory and are named "Abscissa." Abscissa files can be multicolumn ASCII files. The abscissa is used for several purposes, for example it is used to generate ASCII files from image pixel data. Additionally, when an image is processed on a pixel by pixel basis, the model that does the processing must know the abscissa values. For example if one is processing diffusion tensor data, one must know the "B" values. The "B" values are the abscissa, and in this example the abscissa file would be a three column ASCII file. The number of columns in the selected Abscissa file must match the requirements of the current package before the interface will load the file. When the Load Abscissa button is activated, one can select one of two options:

From File will bring up a popup that allows navigating to and loading of an Abscissa file.

From Procpar will bring up a popup that allows navigating to and selecting a procpar file. The interface will then read the procpar and find the arrayed variable and attempt to construct the abscissa from procpar. If multiple variables are arrayed in the procpar, then a multicolumn Abscissa file is constructed.

Download Test Data to the Bayes Home directory. The Bayesian Analysis software ships with a directory containing data that can be used to test the various packages. This data is contained in a file on the server. That file is located in the Bayes user account in a directory named Bayes. The file is named Bayes.test.data.tar.gz. However, the data is not generally accessible to users because the file is a gzip compressed tar file. When selected, the "Download Test Data" submenu, downloads a copy of this file and then uncompress and untar the files. The directory, Bayes.test.data, containing the test data is placed in the Bayes directory in the bayes user account. The downloaded test data, can then be load and used to test the various packages. Inside the Bayes.test.data directory the test data is organized by package.

Download Manual (pdf) will download a copy of the user manual to the Bayes Subdirectory in the current Bayes Home directory. This manual is named BayesManual.pdf and is the version of the manual issued with the Bayesian Analysis software installed on the server. Additionally, a web browser can be pointed to

http://bayes.wustl.edu/Manual/BayesManual.pdf

and the most recent copy of the manual can be downloaded from the Bayesian Analysis' home page. Finally, the above link can be activated and the acrobat reader will download the most recent version of the manual from bayes.wustl.edu.

Save Working Directory widget allows a WorkDir to be saved. When this widget is activated, a popup is displayed. This popup allows selection of both the WorkDir and the location where the WorkDir is to be saved. Finally, the WorkDir is copied to the specified location. When the WorkDir is copied, the entire contents of the WorkDir are copied, all ASCII files, images, fid's etc. are copied and saved in the specified location.

Import Working Directory allows a saved WorkDir to be reloaded. When the Load Working Directory widget is activated a popup is displayed that allows the select the WorkDir to be reloaded. In a copy statement this is the source location of the files to be copied. After selecting the source WorkDir, you are prompted to enter the name of the WorkDir where the files are to be copied will appear. The source directory is then copied to the Bayes directory in the current Bayes Home using the new WorkDir name as the "to" location in the copy. After reloading the WorkDir, plots, text reports, fid's, images and ASCII files will have been restored to the same status they were in when the WorkDir was saved.

Import Working Directories in Batch imports multiple saved working directories. For example while working on a project where model selection was needed on about 100 different samples. We created a working directory for each sample, ran the model selection and then save the resulting working directories. At a latter time it was necessary to reload all 100 working directories so that we could review and in some cases rerun the analysis.

3.1.2 the Packages menu

The packages menu, shown in Fig. 3.4, is used to select a package, a set of programs used to solve some particular problem. The software contains roughly 20 packages and these packages implement various calculations using Bayesian probability theory. The various packages implemented by this software are briefly describe here and a detail description of each package is given in later Chapters.

Exponential The Exponential package estimates the decay rate constants and amplitudes of signals known to be decaying exponentially. It does this when the number of the exponentials is known or unknown. In both cases the input to this package can come from ASCII files, from a peak pick or from Bayes Analyze files. In all cases one or more input data sets can be processed. When multiple input data sets are processed, the package looks for exponentially decaying signals that are common to the various data sets, but allowing each exponential to have differing initial conditions in each data set. See Chapter 5 for more on the exponential problem when the model is given, and Chapter 6 when the number of exponentials or the present of a constant offset are unknown.

Inversion Recovery The Inversion Recovery package is a special type of exponential analysis that is very common in NMR. In this problem the NMR signal starts at a negative value and decays to a positive value. The inversion recover model differs from an exponential plus a constant model only in that the model is typically formulated so that the two amplitudes represent the initial, time equal to zero, and equilibrium amplitude; thus the amplitudes are linear combinations of the amplitudes that would be estimated by an exponential plus a constant model. As a side note, this package is really a special case of the Enter ASCII package described below. We call these special cases preloaded enter ASCII models because the interface preloads the inversion recover model from the system model directory and thus simplifies what the user must do to run this inversion recovery model. This package can analyze multiple data jointly to look for a common parameters. See Chapter 7 for a description of the inversion recovery package.

Diffusion Tensor The Diffusion Tensor package analyzes NMR diffusion measurements using one, two or three diffusion tensor models with or without a constant. These tensor can use either "b"

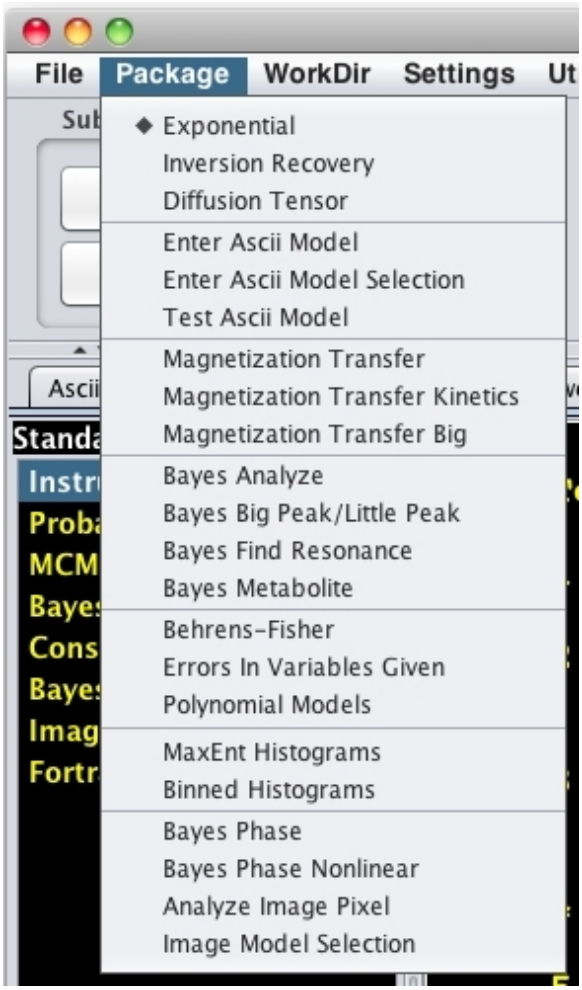

The package menu allows selection the of the package to be run. When activated the menu lists all available packages. The packages are grouped more or less by the type of data and model processed. For example, Exponential, Inversion Recovery and Diffusion Tensor all process ASCII data and they all process models that are exponential in nature; while things like Bayes Analyze, Big Peak/Little Peak, Find Resonance, and Metabolite analysis all analyze fid data and they all estimate parameters associated with resonances. A brief description of each package is given in this section and a Chapter is devoted to describing the models and in some cases the Bayesian calculations done in each package.

Figure 3.4: When the Package menu is selected this pull down menu is displayed. It is populated with a complete list of all of the packages supported by the Bayesian Analysis software. Selecting a package will cause the interface to display the interface to the selected package and the interface will configure the current working directory for that package.

values or "g" (gradient) values for the abscissa and the "b" values can be either 3D vectors or "b" matrices. Thus this package process 18 different diffusion tensor models. Because McMC packages compute the probability for the model using thermodynamic integration, this package has the ability to do some simple model selection. As with most packages multiple ASCII data sets can be analyzed jointly to look for common diffusion tensor parameters. process ASCII Diffusion tensor models, similarly for image model selection. See Chapter 12 for a description of the diffusion tensor model.

Enter ASCII Model The Enter ASCII Model package allows the user to define a model and then use Bayesian Probability theory to analyze data using that model. To create a simple model, activate the Fortran/C Code Viewer and then activate the "Edit/Create New Model" button. When this button is activated it will make a copy of the Example.f model, and open it in an editor. This model can be changed, compiled, tested, save and run as needed. In addition to creating a Fortran/C model, the users must create a file that describes the model parameters and the prior probabilities for those parameters. This process is done simiautomatically when the Fortran/C model editor is used. However, If models are edited manually then this file must be created manually. See Chapter 20 for a description of the Fortran/C models and their "params" file.

Enter ASCII Model Selection The Enter ASCII Model Selection package utilizes the models generated for Enter ASCII to do model selection. After setting up a number of rival models using Enter ASCII, one can then proceed to this package. Here one can load up to 10 different models and then use this package to compute the posterior probability for the models. The only requirement between the models, is that they must process the same data, so all models must have the same number of data columns, and because ASCII data has the abscissa in the file all models must use the same abscissa. See Chapter 22 for a description of the ASCII Model Selection package.

Test ASCII Model The Test ASCII Model model package supports the other packages that use ASCII Models. This package gives one a facility for testing models to ensure they are doing their calculations correctly. This package allows load a model and data associated with that model, and then the Test ASCII Model package will thoroughly test the model by evaluating the model 10,000 times using parameter sampled from the priors. In the process of evaluating the model, the package will catch any arithmetic errors that occur and it will show the abscissa value where the invalid arithmetic occurred. The outputs form the model include a peak posterior probability estimate of the model and plots of the model signal as a function of the parameter samples and plots of the residuals, the difference between the data and the model. See Chapter 21 for a description of the test ASCII model package.

Magnetization Transfer The Magnetization Transfer (two sites) package solves the Block-McConnell equations to obtain the exchange rate constants for two site magnetization exchange. Input to this package is usually the peak amplitudes or intensities from two inversion recovery time courses where the exchanging peaks in are selectively inverted. The ASCII file used by this package is three column ASCII, one abscissa and the amplitudes of the two exchanging peaks. See Chapter 14 for a description of the Magnetization Transfer package.

Magnetization Transfer Kinetics The Magnetization Transfer Kinetics package is a magnetization transfer package that solves the Block-McConnell equations at multiple temperatures and

concentrations to derive the entropy and enthalpies of the the exchange process. Input to this package is also three column ASCII, with multiple data sets taken at differing temperature and concentrations. See Chapter 15 for a description of the Magnetization Transfer Kinetics package.

Big Magnetization Transfer The Big Magnetization Transfer package solves the magnetization transfer problem when one of the sites can be considered infinite compared to the other. See Chapter 13 for a description of the Big Magnetization Transfer package.

Bayes Analyze The Bayes Analyze package is a time domain frequency estimation package that is fully capable of determining the number of resonances in an fid and estimating the resonance parameters. This package can analyze single fid's, or it can run multiple fid's and look for frequencies common to these fid's. Input to this package can come from different sources and appropriate data conversions are carried out when the data are loaded. See Chapter 8.1 for a description of the Bayes Analyze package.

Big Peak/Little Peak The Big Peak/Little Peak package analyzes time domain fid data in which there is a single big peak that may be many orders of magnitude larger in intensity (the big peak) than the metabolic peaks (the little peaks) of interest. The Big Peak/Little Peak package solves this problem by treating the big peak as a nuisance and then uses Bayesian probability theory to account for the big peak while simultaneously estimating the frequencies, decay rate constants and amplitudes of the resonances of interest. See Chapter 9 for a description of the Big Peak/Little Peak package.

Find Resonances The Find Resonances package analyzes NMR fid data looking for resonances. The program is a model selection program that is attempting to determine the number of resonances in the data and estimate the parameters associated with those resonances. This package uses Markov chain Monte Carlo simulations to determine the posterior probability for the number of resonances in the data. This package essentially solves the same problem as the Bayes Analyze package described above. However, because it uses McMC the calculations are much slower than those in Bayes Analyze, but they are much more through; often having much better resolution than Bayes Analyze. See Chapter 11 for a description of the Find Resonance package.

Metabolite The Metabolite package analyzes fid data from a number of known samples, for example a C13 fid of Glutamate. The intensity of the Glutamate resonances are related to each other through a metabolic model. This model can be very simple or very complex. Metabolic models can be added to the library of models, but there are no facilities for building these models within the interface. Metabolic models relate the intensity of the resonances in the model to a series of metabolic parameters, typically fractional rates that relates how much of a compound went through a certain chemical reaction. The resonances in a metabolic models are described in a metabolite file and the metabolic model itself is encoded in a FORTRAN or C routine. The metabolic package reads the resonance and the metabolic models and then uses Bayesian probability theory to estimate the metabolic parameters as well as the parameters associated with the resonances, i.e., the frequencies and decay rate constants. See Chapter 10 for a description of the Metabolite package.

Behrens-Fisher The Behrens-Fisher package solves the classical medical testing problem: given two experiments that consist of repeated measurements of the same quantity where in the second measurement one has change some experiential parameter determine if the experiments are the same or if they differ. See Chapter 19 for a description of the Behrens-Fisher package and see On the Difference in Means for a detailed description of the calculations.

Errors in Variables The Errors in Variables package solves the errors in variables problem. In this problem one has a data set that has uncertainty in both the X and Y variables. These errors may be know or unknown, so this package solves four different errors in variables problems. In the name the "given" refers to the fact that the program solves this problem given the order of the polynomial to fit. See Chapter 18 for a description of the Errors in Variables package.

Polynomial The Polynomial Models package fits polynomials of either a given or an unknown order to the input data. When the order is specified then a polynomial of that order is analyzed using Bayesian probability theory to determine the appropriate coefficients. When the order is specified as unknown, the Bayesian probability theory is used to compute the posterior probability for the order of the polynomials. The input data is two column ASCII and this package do not process multiple data sets. See Chapter 16 for a description of the Polynomial package when the order of the polynomial is given and see Chapter 17 for the calculations when the order of the polynomial is unknown.

MaxEnt Histograms The Maximum Entropy Histograms, density estimation package, is a ASCII package that takes as its input a two column ASCII file. Column one is just a data point number and column two is a sample from the unknown density function. The program models the density function as a Maximum Entropy moment distribution having an unknown number of Lagrange multipliers. So the parameters are Lagrange multipliers and the unknown number of them. The program does a Markov chain Monte Carlo simulation with simulated annealing where the number of multipliers is one more parameter in the simulation. Outputs include the posterior probability for the number of multipliers, the posterior probabilities for the multipliers, scatter plots and the polynomials used in the calculations. See Chapter **??** for a description of the Maximum Entropy Histograms package.

Binned Histograms The Binned Histogram package is a new histogram package. In the previous release of the software, there was a MaxEnt histogram package that infers histograms that are functionally Maximum Entropy moment distributions. As such the program is inferring the moments and the number of moments needed to represent the input samples from unknown density. This procedure works well for compact distribution, but fails badly when the distribution of samples is multimodal. In order to estimate density functions when the samples are multimodal we added a histogram package that infers what can only be called binned histograms. These histograms can represent any distribution, they have error bars on the number of counts in the bins, and the user can indicate if the histograms are to be smoothed or not. See Chapter **??** for a description of the Binned Histograms package.

Linear Phasing The Linear Phasing package produces linearly phased images. In spin echo MRI most images can be phased (absorption mode images) by calculating two first order phases and one zero order phase. Bayes Phase computes these phases and then applies them to the images. The resulting images are then available for further processing by the Analyze Image

Pixels package. See Automatic phasing of MR images. Part I: Linearly varying phase for more on this calculation and and for a more detailed description of this package see Chapter 23.

Non-Linear Image Phasing The Non-Linear phasing package phases images that are varying in a Non-Linear fashion. This package takes as its input the output from the Linear Phasing package. This package can be used to produce absorption mode images for gradient echo MR images or any other image in which the phase is varying in an unpredictable fashion. For more information on this calculation see Automatic phasing of MR images. Part II: Voxel-wise phase estimation and for a more detailed description of this package see Chapter 24.

Image Pixels The Image Pixels package loads a predefined model and then uses that model to analyze images on a pixel by pixel basis. Model can be loaded from the system directory and these predefined models perform a number of common calculations in MRI such as exponential analysis with one or more exponentials with or without a constant, diffusion tensor, Additionally, the users can copy and the edit an example model to create models of his own. These models can be loaded from the users home directory and then used to analyze the image.

The Image Pixels package includes an option for finding the peak of the posterior probability. When this option is selected, a different program is actually run by the package. This program is a searching algorithm that looks for the peak in the posterior probability for the parameters in the model. These peak parameter estimates are then used to generate maps of the various parameters appearing in the model. Because this program is a searching routine rather than an McMC routine, it is very fast and can give good results using any ASCII model in a fraction of the time needed to run the Markov chain Monte Carlo simulations. See Chapter 28 for a description of the Image Pixel package.

Image Pixel Model Selection The Image Pixels Model Selection package extends the concepts in Analyze Image Pixels to model selection. In this package one can load a number of different models and then use Bayesian probability theory to determine which model best accounts for the data. The models in use here are the same models mentioned in both Analyze Image Pixels and the Enter ASCII packages. However, here because the models can have different parameterizations, the output images are constructed from the derived parameters. See Chapter 29 for a description of the Image Pixel Model Selection package.

3.1.3 the WorkDir menu

Working directories are directories that are used to run, configure and store analysis in the Bayesian Analysis Software package. They are physically located in the current Bayes Home directory. The default location of the Bayes Home directory is the user home directory and the default name of the Bayes Home directory is “Bayes”. The name and location of the Bayes Home directory can be configured using the Settings/preferences popup and multiple Home directories are allowed. The working directory menu, called WorkDir, is generated on the fly and contains a list of the working directories in the current Bayes Home directory. Additionally, the WorkDir menu contains one fixed menu item named “Edit”. The WorkDir menu is a pull down menu that allows the user to manage working directories, Fig. 3.1.3. The top part of this menu will list all of the working directories in the current Bayes Home directory. The name of the working directories are assigned by the user and the name can be descriptive of the type of analysis being done in that directory. For example, in Fig. 3.1.3 all of the working directories have names that indicate of project the user was working on.

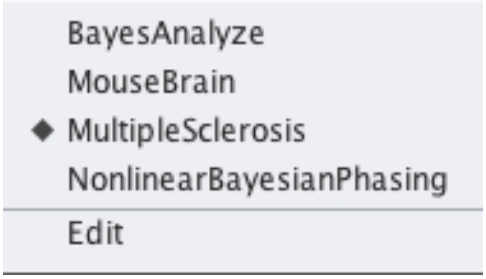

The contents of this menu varies depending on the number of working directories and the names that have assigned. By selecting one of menu items the interface will join that working directory and if an analysis is present it will restore that analysis to its previous status.

Figure 3.5: When the WorkDir menu is selected this pull down menu is displayed. It contains a list of all of the current Working directories.

When a working directory is selected, the interface will save the status of the current WorkDir and then change into the selected working directory and restore that working directory to its previous status.

The last entry on this menu is the "Edit" button. When activated, the edit button will bring up the popup shown in Fig. 3.6. that allows one to modify and manage working directories. Along the left-hand side of this menu is a list of all of the working directories. As noted, these are the names of the working directories as defined by the user. By clicking on these items, the status of the working directory is displayed. Status information includes, the package that is loaded, whether or not the package has been run, and information about the files, ASCII, fid, and Images that have been loaded. Additionally, information about the server that is selected in this WorkDir is displayed. Finally, using the buttons along the button of the WorkDir manager, a new WorkDir can be created. To create a WorkDir, simply enter the name of the directory in the text area in front of the "New" button. and click the new button to create the directory. To delete a WorkDir, select the directory to be delete, and activate the delete directory button. Finally, to load a WorkDir, select the directory to be loaded, and activate the load directory button.

3.1.4 the Settings menu

The settings menu is shown in Fig. 3.7. This menu allows one to configure the interface to make use of the operating environment and to control the Markov chain Monte Carlo simulations. Set the window size and to set so user preferences. Here is a description of the this Settings menu:

the McMC Parameters submenu Many, indeed, most of the Bayesian Analysis packages use Markov chains to approximate the joint posterior probability for the parameters appearing in the model. This is done by using the Markov chain to draw samples from the joint posterior probability for of the parameters. From these samples, Monte Carlo integration can be used to approximate the posterior probability for each parameter appearing in the model. The number of Markov chain, the number of samples gathered and the annealing schedule are things that the user can control, i.e., configure. Activating the McMC parameters widget brings up the popup shown in Fig. 3.8. This popup allows number of Markov chain Monte Carlo simulations that are to be run concurrently or in parallel to be set: Concurrently if only a single processor is available, and in parallel if multiple processors are available. Additionally, the number of

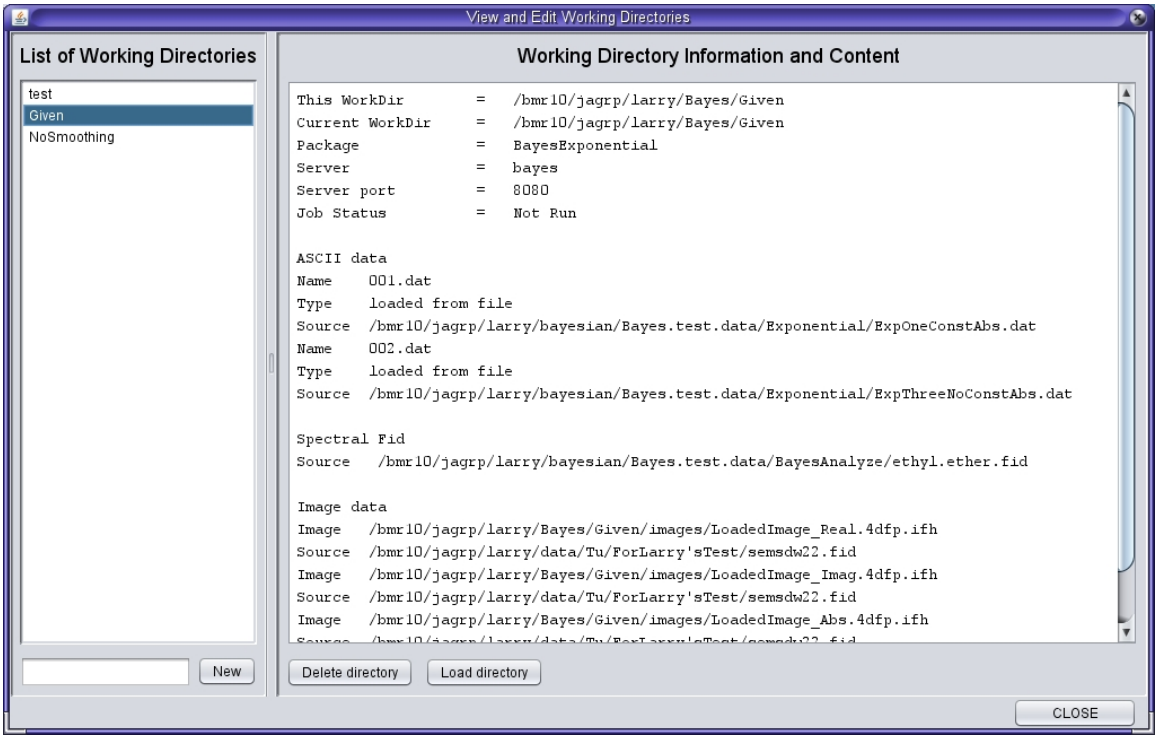

Figure 3.6: When the WorkDir Manager is selected this popup window is displayed. Along the left-hand side is the list of the working directories. By clicking on these working directories, the current status of each directory can be viewed and by using the buttons at the bottom working directories can be created, loaded or deleted.

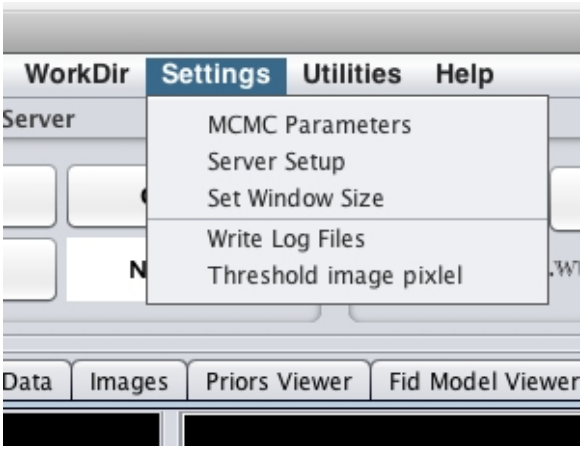

The Settings menu is a pull down that allows one to configure and control a number of important features of the interface. The two most important settings concern the McMC parameters and the Server Setup. When the McMC Parameters menu is activated, it brings up a popup that allows the number of simulations, repeats and the minimum number of annealing steps used in the Markov chain Monte Carlo simulations to be set. When the Server Setup is activated, the resulting popup allows one to add, remove and configure servers. Finally, the preferences button will bring up a popup that allows the configuration some interface parameters, like for example the location of the Bayes Home directory.

Figure 3.7: When the Settings menu is selected this pull down menu is displayed. It allows one to configure the Markov chain Monte Carlo simulations and to configure a series of users preferences.

The McMC Settings menu is a popup that set the number of Markov chain Monte Carlo simulations that run concurrently or in parallel: Concurrently if only a single processor is available, and in parallel if multiple processors are available. Additionally, the number of McMC repetitions can be set, and thus the number or samples gathered for use in computing mean and standard deviations parameter estimates. The number of samples = number of repeats times number of simulations. Finally, the minimum number of annealing steps to take during the the simulated annealing phase can be set. For more on how the McMC is used in the Bayesian Analysis software, see Section B.

Figure 3.8: When the McMC Parameters menu is selected this popup is displayed. It allows a number of important parameters concerning the Markov chain Monte Carlo simulations to be configured.

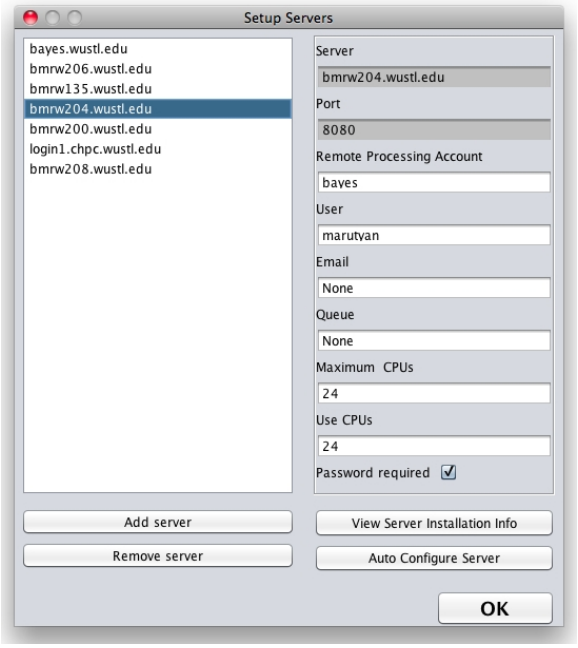

The Server/Server Setup menu is a popup that allows one to add, delete and modify server settings. To select a server simply click on the server name. Any field except the server name and port number can be modified. Servers can be added using the Add Server button. Servers can removed by activating the remove server button. Finally, the View Server Installation Info button will bring up a popup that lists all of the installation information available on the server.

Figure 3.9: When the Server Setup menu is selected this popup is displayed. It allows Servers to be add, remove and change.

McMC repetitions can be set, and thus the number or samples gathered for use in computing mean and standard deviations parameter estimates. The number of samples = number of repeats times number of simulations. Finally, the minimum number of annealing steps to take during the the simulated annealing phase can be set. For more on how the McMC simulations are run see Section B.

the Server Setup submenu Activating, the Settings/Server Setup menu will bring up the popup shown in Fig. 3.9. The Settings "Server Setup" menu is a popup that allows servers to be add, delete and modify. To select a server simply click on the server name. After selecting a server, that server becomes the current server and any job submitted will be sent to the selected server. Note that servers can also be selected by activating Server/Set button on all package interfaces. When servers are selected on the Settings/Server Setup popup, the server name, port, Bayes user account, etc. are displayed. Any field except the server name and port number can be modified. However, modifying anything other then the user name and email preferences is not advised. Indeed, modifying the processing account, password settings or queue name will likely result in nonfunctional server configuration.

The Settings "Server Setup" popup has four buttons that can be used to add, delete, configure and display server information.

Add Server will bring up a popup in which the server name and port number of the server can be entered. After entering the server name and port, most of the server information

can be configured using the Auto Config Server button. However, email preferences and user name may have to be set manually.

Remove Server will delete the currently selected server from the list of servers. Note there is no prompt to ask if a server is to be removed.

View Server Installation Info will bring up a popup that shows more information about the server. In will show the date the software was installed, the software version, the Bayes user account, port, compiler information including the path to the compilers, number of CPU, whether or not passwords are in use and information about who the Bayesian Analysis administrator is.

Auto Configure Server fetches a configuration file from the server. It then uses this configuration to set the processing account, passwords and number of CPU on the Server Setup popup.

Set Window Size will bring up a popup that allows the size of the interface windows to be set. Normally one would set the size of the interface by using a cursor to stretch/contract the window as needed. However, sometimes use of a cursor is difficult depending on the type of terminal in use. This menu allows the size of the window to be set without using a mouse.

Preferences The preference widget will bring up a popup that allows you to configure some preferences. The main settings are for the location of the Bayes Home directory. You can uninstall the Java interface from you server. You can tell the interface what output format to write screen captures in. Finally, there are a couple of widgets that indicate if jobs are to be deleted from the servers. This last function is mostly used by us in debugging software. It could happen that while trying to diagnosis a problem, we need for the job to remain on the server so we can see what happened to it. Normally, completed jobs, i.e., any job that is no longer active on the server are removed as soon as the get job widget is activated.

3.1.5 the Utilities menu

The utilities menu allows the user to run a number of utilities. These utilities can monitor the memory usage, display some information about Java and its instillation and finally determine if a new version of the software is available. Here is a more detailed description of these utilities: There are three utilities that can be run by the user:

Memory Monitor will activate a Java Memory monitor. Java, for whatever reason, normally can only access 1GB of memory. Sometimes, we have found that this is not enough to hold large images and applications. The memory monitor can be used to monitor memory usage. If 1GB of memory is not enough, it is possible to get 2GB of memory. However, the launch.jnlp files on the server must be modified to do this. If assistance is needed in modifying these files, please contact.

System Info will display information about the current installation on the client machine. This information could be important to us when trying to help diagnosis problems. The displayed information includes: the user name, the user home directory, the architecture, OS, Os version, file separator and the version of Java in use. In case of problems all of this information could be useful in figuring out what is going on.

Software Update will bring up a popup showing all of the servers currently have defined and it will indicate if each server has the most recent version of the software installed on it. So the Software Update utility check with us here at Washington University to find out what the current version of the software is and it then checks each of the servers to see if they are up-to-date. If the software is not up-to-date log into each out-of-date server as the bayes user and then update, i.e. reinstall, the software.

3.1.6 the Help menu

The Help menu, not shown, will provide help concerning the current release of the software. There are four different types of help available:

Release Notes will activate a web browser and download a page that describes the changes in the current version of the Bayesian Analysis Software. Those release notes contain links to the main BayesianAnalysis home page and a chain of links that will describe in the previous releases.

Online Manual will download the current version of the manual and the display that manual in the Acrobat reader.

Bayes Analysis Home Page will load the home page from the BayesianAnalysis.wustl.edu web site and bring that page up in the default web browser. That page contains a description of the software as well as the release notes for the current and previous versions of the software.

Contact Us bring up an email client with my email address in it so that questions can be directly addressed to me.

3.2 The Submit Job To Server area

Just below the global pull down menus, there are a number of widget groups that are used to configure a package. These widget groups are different for different packages. And details on a package specific widget group, consult the chapter on that package. However, their are two widget groups that are global in the sense that they occur on all packages. In the next two Subsections we are going to describe these widget groups and their function.

The first of these widget groups is called the "Submit Job To Server" widget group and this widget group is shown in Fig. 3.10. Below the global menus is an area that is used to configure a package.

Run will first check to see if all the required elements are set for this particular job. If everything is properly configured it will create a tar file of the current WorkDir and send it to the server shown in the server widget group. When the job arrives at the server, it is untared and then and the requested programs are run. After the job arrives at the server, interface set the status to either Active or Queued and it will lock most of the widgets on a package. The user is free to join another analysis or wait for the analysis to finish.

Cancel will cancel a submitted job and remove all files pertaining to that job from the server even if the job is completed. When finished canceling the job, the interface will set the current

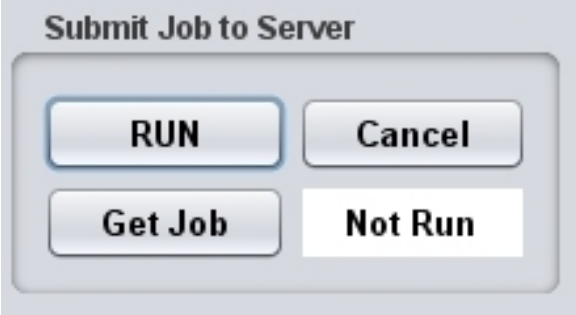

Run will submit a job to the indicated server.

Cancel will cancel a submitted job and remove all files pertaining to a job.

Get Job will fetch the current job or fetch the status of the current job.

Status Label contains the status of the current job.

Figure 3.10: The submit Job to the Server Widget Group. Is used to send and fetch jobs from the servers.

status to either "Not Run" or "error" if the job could not be canceled and or removed from the server.

Get Job performs two main tasks, when activated it first checks the current status of a job and depending on the status it either fetches the job or it fetches and displays the accepted report.

Text Label contains the status of the current job. Note this field is only updated any when one of the three Submit Job to Server buttons are activated. To obtain the current status of a job, activate the "Get Job" button.

3.3 The Server area

The "Server" widgets group configures and controls server. This widget group is shown in Fig. 3.11. In general terms this widget group allows the select and configuration of servers. The server widgets group

- The server "Set" button allows the current server to be selected. When this button is activated, a pull down menu appears containing a list of all of the servers. Clicking on a server, will cause it to be set as the current server. The current server is displayed in the server name text are under this button. At the bottom of pull down menu is an item "Edit Servers" that can be used to modify the list of servers. Activating this widget will bring up a popup, Chapter 3.1.4, that allows servers to be add, deleted and modified as desired. This Server Edit popup is also available under the "Settings/Server Setup" menu.
- The server "Status" button will send a request for a list of jobs currently running on the server. On Linux and Sun systems this request is a simple "ps". The results of this request are displayed in the Text Viewer at the bottom of the interface.
- The current Server is displayed in the "Server Name" text area under the two button in the Server widget group.

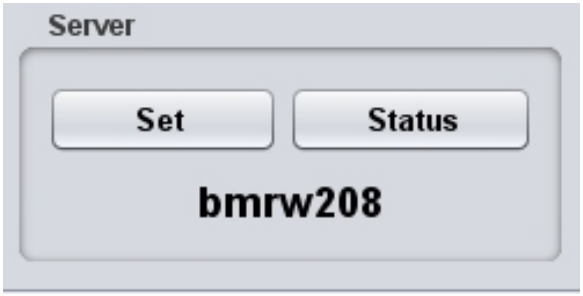

Set will bring up a popup window containing a selection list of the current servers. Selecting one of these servers will change the current server. The Server Name will be changed to reflect the selection. Note, selecting the "Edit" button in the server list will bring up the Server Configuration popup discussed in Subsection 3.1.4.

Status will send a request to the selected server asking for its current CPU load. This system load is displayed using "ps" on most Linux systems.

Server Name contains the name of the currently selected server.

Figure 3.11: The Server Widget Group is used to setup, change or to check the status of a server.

3.4 Interface Viewers

Just below the widget groups is a set of buttons that activate various viewers. These buttons start on the left in Fig. 1.2 with Ascii Data Viewer and end on the right with File Data Viewer. Each viewer is used to look at a given type of data. On the Exponential package there are seven of these viewers, and this is pretty typical of all packages.

3.4.1 the Ascii Data Viewer

Ascii data can be loaded and viewed in all packages, even packages that analyze fid and image data. The Ascii Data Viewer, shown in Fig. 3.12, is used to display this data. To load an Ascii data select the "File/Load Ascii/File" submenu item. When activated this widget will bring up a popup file loading widget. Navigate to the desired file and then select and load the file. Please note that there are some rather complicated rules concerning what Ascii files can be loaded in a given package. To give one example in the exponential package only load two-column Ascii files can be loaded, while in the magnetization transfer packages only load three column ASCII file can be loaded. Additionally, files suitable for the the magnetization transfer packages will not load in the exponential packages and vice versa, i.e., the file loading popup knows how many abscissa and data columns are required for a given package/user define model. For the exact requirements for each package, see the Chapter on that package. When an Ascii data set is loaded the data is copied into the BayesOtherAnalysis directory of the current WorkDir and the data set is renamed as 001.dat, 002.dat, etc. where the assigned number is just the number of the loaded Ascii data set. The data sets are renamed with unique names to prevent name collision problems. The original name, size, number of columns are stored in a file named 001.afh, 002.aft, etc. These Ascii file header files, the "afh" files, are used to display information about a Ascii file. When an Ascii data set is successfully loaded the Ascii Data Viewer, Fig. 3.12, is automatically activated. Additionally, it can be activated by clicking the Ascii Data Viewer button. When activated a list of all loaded Ascii data files is shown on the left-hand side of the viewer. When a left-hand file name is selected by clicking on it, the selected data set is

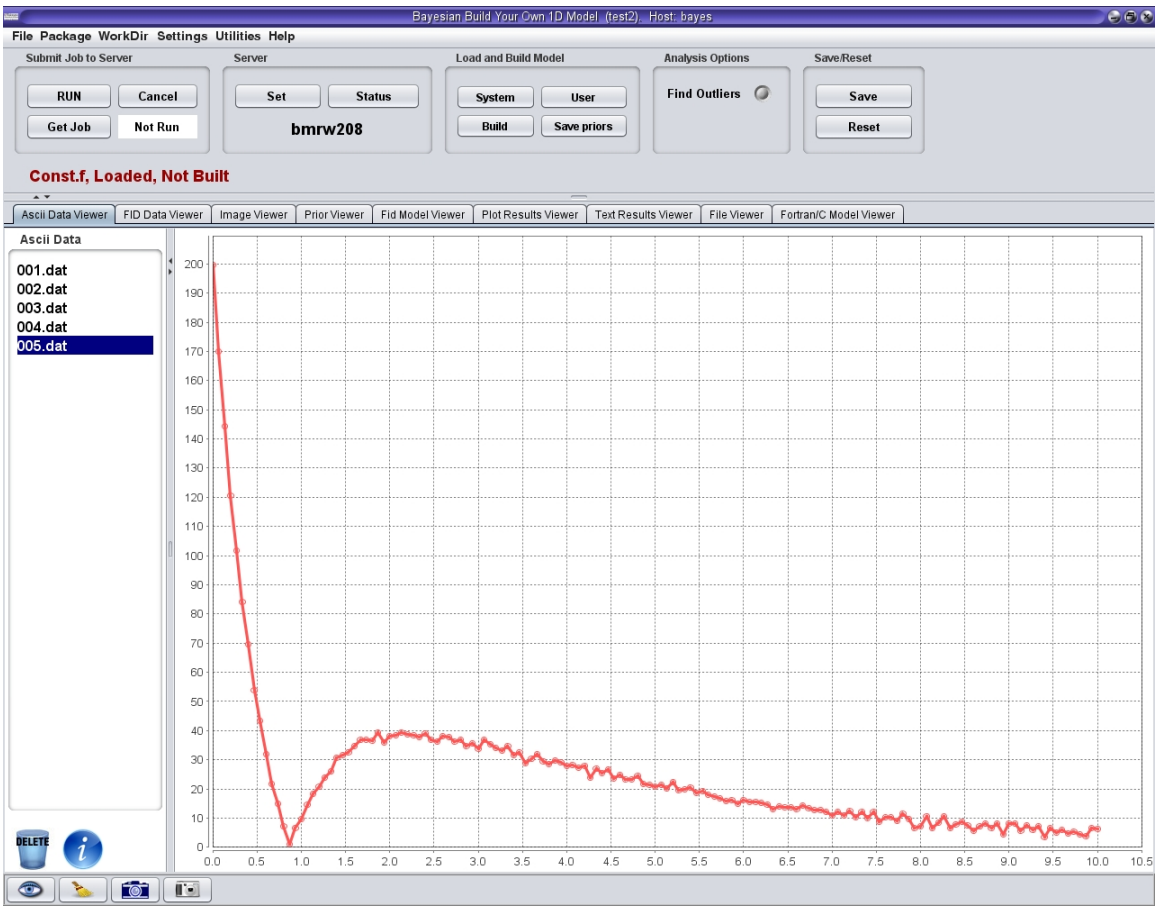

Figure 3.12: The Ascii Data Viewer is used to display Ascii files. Change files by clicking on the name, of the file to be viewed. The "Delete" button will delete the selected file and the italics "*i*" button in the circle will display everything known about the file. This button is redundant with a right mouse and selecting show info.

plotted on the right-hand side viewer. If the data are multicolumn data, for example complex data, the plot will have multiple traces on it.

This viewer responds to a right-mouse click in both parts of the viewer. In the file list area, the left-hand side, the right mouse click shows a submenu that allows file deletion, displays information about the selected file and allows one to view the Ascii file as text. In the plotting area, a right mouse click brings up a submenu that allows one to configure the plot. For example a right-mouse click to can be used to change the headings, axis labels and to save or print the plot.

On the left-hand side at the bottom there are two widgets, one labeled "Delete" and one labeled "*i*" that delete the selected file or show information about the selected file. These functions are redundant with the functions available using a right-mouse click.

Finally, on the plotted area, a left-mouse click and dragging the cursor down and to the right will select, highlight, a region. When the cursor is released, the highlighted are is expanded. To return to the full plot, left-mouse click and drag the cursor up and to the left will restore the plot.

3.4.2 the fid Data Viewer

The Fid Data Viewer, shown in Fig. 3.13, is activated whenever fid data is loaded from the Files menu, or when the Fid Data Viewer is selected. This viewer allows one to look at both the time and frequency domain fid data. When an fid is loaded, the fid is copied into the "fid" directory in the current WorkDir and written in Varian format. The data is then Fourier transformed, the Fid Data Viewer is activated, the spectral data is phased and plotted. Additionally, like the "afh" file an "ffh" file is written that contains information about the fid. This "ffh" file is displayed whenever information about the currently loaded fid is requested.

The Fid Data Viewer uses left, right and center mouse clicks. The left and right mouse clicks are used to set the locations of a left and right cursor. These cursors are displayed in Fig 3.13 as the two vertical red lines. When set the frequencies of these cursors are shown in in at the bottom of the viewer as "Cursor A" the left cursor and "Cursor B" the right cursor. The difference between these cursors is shown in the "Delta" display. Note the units used in these three display areas are the same as the units on the displayed frequency axis. When the left and right cursors are displayed, the "Expand" button can be used to expand the selected area. The "Full" button will display the full spectrum and the "Clear Cursors" will remove the cursors from the display.

The vertical scale on the display can be adjusted automatically using the "Autoscale" button or it can be adjusted manually. The center mouse button is used to adjust the vertical scale manually. If cursor is placed anywhere in spectrum display above the axis and hit the center mouse button, the display vertical scale is adjusted upward. If cursor is placed below the axis and hit the center mouse button the vertical scale is reduced.

In addition to adjusting the vertical scale, the position of the spectrum in the viewer can be changed. The left-hand part of this viewer has a vertical strip that is shaded a slightly different color than the black of the viewing area. This area can be used to adjust the position of the display vertically. So if the base-line of the spectrum is too high or too low place the cursor in this shaded region and then hold down the center mouse button and the display can be moved up or down.

The type of display can be changed on the Fid Data Viewer. The top-center pull down menu that usually reads "SPECTRUM REAL" can be used to change the type of display. When this pull down menu is extended the menu shown in Fig. 3.14 is displayed. Most of the options are self explanatory and we give no further explanations of this viewer here except to note that all of the cursor and mouse functions work on the different data types.

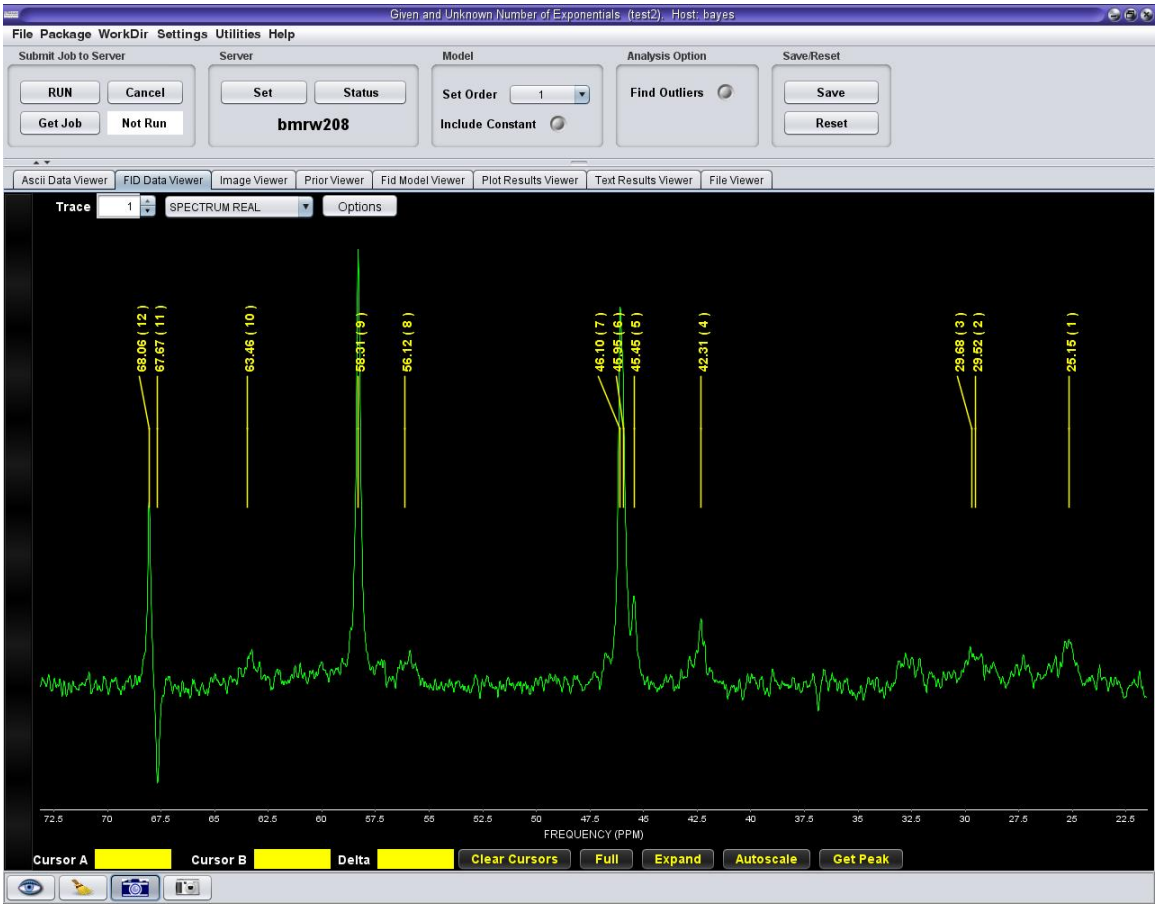

Figure 3.13: The Fid Data Viewer is used to display fid files. When activated the spectrum of the currently loaded fid is displayed. If no fid data are loaded an empty viewer is display. This viewer can be used to zoom in on resonances, change the scale, phase and many other functions. See the text for an extensive discussion of this viewer.

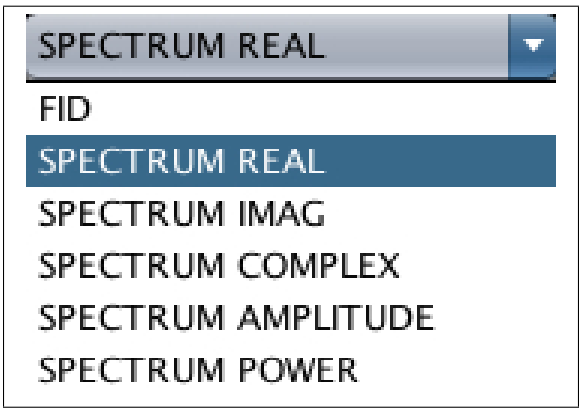

This pull down menu allows selection of type of data to be displayed. The options allow the display the fid or the spectrum in a number of different modes such as absolute value, power spectrum, etc. The real or absorption mode spectrum is the default. Note that this widget determines the type of data displayed, but it does not determine the form of the report.

Figure 3.14: When the Data Type menu is selected this pull down menu is displayed. This menu sets the type of data that is to be output. For example, selecting "fid" will cause the time-domain fid data to be displayed.

If a peak is expanded and the "Get Peak" button is activated, the amplitudes of the peaks will be extracted from the current cursor position. The interface will attempt to combine these peak amplitudes with any arrayed variable from the fid procpar to produce an Ascii data set that has the appropriate axis. This data set is assigned a new Ascii data set number, and the Ascii Data Viewer is activated to display the data.

The Options widget is the most complicated widget on the Fid Data Viewer. Its functions are extremely varied and describing all of the functions hidden under the "Options" menu would be very difficult. Here we are going to list these options and give some pointers on what they do. First the expanded options menu is shown in Fig. 3.15. Here is a brief explanation of the various options:

Data Info Brings up a popup window that displays information about the currently loaded fid. This information includes things like the original source fid name. The current weighting, the units in use, the phase, the reference and a number of other data items. Its basically a dump of all of the relevant parameters concerning this fid.

Save As Varian fid this functions saves the currently loaded fid as a Varian binary fid file and it does this regardless of the input format of the data. So an input data set could have been an Ascii fid and this save command will save the data as a Varian fid. Effectively translating the Ascii data to an fid.

Save As Text Saves the currently loaded fid as a two column Text file: (Real, Imaginary).

Show Plotted Data brings up the current plot in a popup window that can be viewed, printed and saved.

Clear Data will remove the current fid, that is to say the files: fid, text and procpar will be removed from the current WorkDir/fid directory.

Data Info will display information bout the fid.

Save As Varian fid will save the current fid in Varian format.

Save As Text will save the current fid in Text format.

Show Plotted Data will popup the data in a Text browser.

Clear Data will delete the current fid.

Apply Phasing will popup a phasing widget.

Set Regions will popup a regions setting widget.

Set Fn will set the FT size.

Set lb will set the line broadening value.

Set Reference will popup a referencing setting widget

Set Units will allow the units to be changed.

Properties allows some plotting preferences to be set.

Copy copies the current image to the clipboard.

Save As popup a save widget.

Print popguns a printing widget.

Zoom In allows the image to be zoomed in.

Zoom Out allows the image to be zoomed out.

Auto Range fits the data to the display window.

Figure 3.15: When the Options menu is selected this pull down menu is displayed. It performs various optional tasks concerning the Fid Data Viewer. For example the size of the Fourier transform, the weighting or reference on the Fourier transform can be set, etc.

Apply Phasing will bring up a popup window that allows the phase of the currently displayed fid to be set. This popup contains two sliders that allow the zero and first order phase to be set. The displayed fid can be zoomed using the procedures described earlier in Section 3.4.1, so while the displayed fid is not initially zoomed before phasing it can zoom as needed. When the phased has been set, hitting the "Phase fid" button will set the phase on all traces in the data. The popup contains a "Trace" button that allows the displayed trace to be changed and it allows the type of data shown to be changed. So for example one could phase the imaginary part of the spectrum if desired. Finally the popup also allows the phase to be "Reset" to the loaded values.

Set Regions Brings up a popup that displays the currently set regions. Regions, low to high frequency intervals, are used in the Bayes Analyze Regions report to calculate the total intensity in a given set of regions. To use this popup simply set a low-high region using the left and right cursors. When a region is set, hit the "Mark" button to add the region to the regions file. Note the "Delete" button can be used to remove a region and the "Close" button will save the regions file. The regions file is save in the current WorkDir in the BayesAnalyzeFiles Subdirectory. This newly created regions file will be used in the regions report the next time Bayes Analyze package is run.

Set Fn is a pull down menu that allows the size of the Fourier transform to be set. The pull down menu contains powers of two from 1K up to 256K that can be used in the calculations.

Set Reference brings up a popup that allows the current reference frequency to be set. At the bottom of this popup are three buttons, "Set Left Most Frequency to Zero", "Set Right Most Frequency to Zero", and "Set Center Frequency to Zero" that are preprogrammed common reference schemes. To set the reference to an arbitrary value, first position the cursor at the point in the spectrum where the reference is to be assigned. Second type into the "New Value" entry box the value of the reference to. Finally, hit the "Set" button to set the reference.

Units is a selection menu that allows the units, "Hertz" or "PPM", to be set.

Properties allows the axis and labels on the Fid Data Viewer to be set. This can be useful when making a graphic that is to be used in a publication.

Copy places a copy of the current graphics window and places it on the clip-board. On Windows and Mac machines this makes the graphics available for plotting and use in papers etc.

Save As will bring up a popup that allows navigation to the directory where the current graphics is to be saved. Enter the name of the file in the "File Name" box. Activate the "Save" button to save a "png" copy of the graphics.

Print will bring up a popup that configures the print jog. After configuring the printer, the "OK" button will print the graphics.

Zoom In will zoom the axis in. There are options for zooming either both axis or either axis separately.

Zoom Out will zoom the axis out. There are options for zooming either both axis or either axis separately.

Auto Range will automatically set the range to view the entire spectrum in both the horizontal and vertical domains.

3.4.3 Image Viewer

Image data can be loaded into any package using the "File/Load Image" menu. As explained in Section 3.1.1, many different types of image can be loaded and these images are converted into 4dfp images and stored in the image Subdirectory. A 4dfp image consists of multiple binary images that are stacked by slice and element number. Multiple images can be loaded provided they have unique names. If the names are not unique then the current image replaces the previously loaded image. Images are used as input to several packages and can be viewed using the Image Viewer. The Image Viewer is is shown in Fig. 3.16. In general terms this viewer consists of four parts: a list of images under the "Image List" label, two sliders that allow a particular image in a stack to be selected.

3.4.3.1 the Image List area

The image list area is used to control what image, a 4dfp stack, is currently being displayed. The image list is just as its name implies, a list of currently loaded 4dfp image stacks. Clicking on a 4dfp stack will cause the stack to be selected and the image indicated by the slice and element number to be displayed. The image within the stack is displayed can be controlled using the sliders at the bottom left of the Image Viewer. The currently selected image stack is highlighted in red and the check box will be checked. Multiple images can be selected either one at a time or in blocks. To select a single image use the control-left mouse click and the image will be selected. Multiple images can be selected by first selecting a single image and then using a shift-left mouse click on another image to select all images between the first and second selected image. Clicking on a single image will deselect all but the single image clicked on. Selected images are used in some of the packages to indicate which images are to be processed by a package.

The Image List and Image Viewing areas responds to a right mouse click. When activated a menu of options appears. this menu is shown in Fig. 3.17. Here is a list of these menu options and the function they perform:

Display Full Image will zoom the display out to show the full image.

Delete Selected will delete the currently displayed image.

Delete All will delete the all images contained in the image Subdirectory.

Autoset Grayscale For Entire Stack will attempt to set a grayscale that can be used do display all images.

Autoset Grayscale For Current Image will attempt to set a grayscale that can be used do display the current images.

View Selected Pixels as Text will display the pixel values contained in an ROI in a popup text window.

Load Selected Pixels will load the pixel values contained in an ROI as an Ascii data set. The abscissa values in this data set are pixel numbers. Loading pixel values will fail if a model is selected that requires more than a two column Ascii Data set.

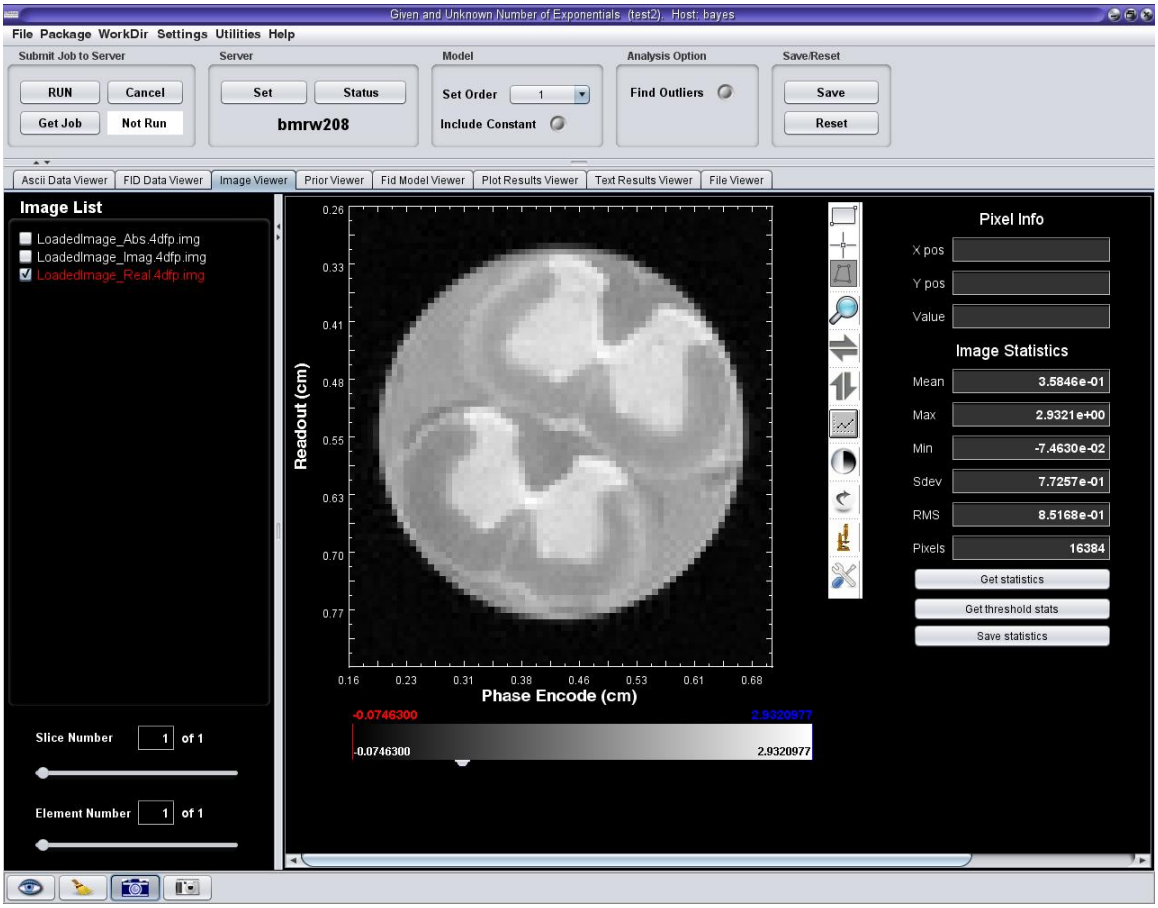

Figure 3.16: When the Image Viewer is selected, this window is displayed. In general terms, this viewer consists of four parts: the image selection widgets on the left-hand side of the viewer. The widgets used to select a slice and element on the lower left-hand side. The image viewing area in the center, and the pixel information area on the right-hand side of the image. For more information on each widget, hold the cursor over a widget and read the tool tip.

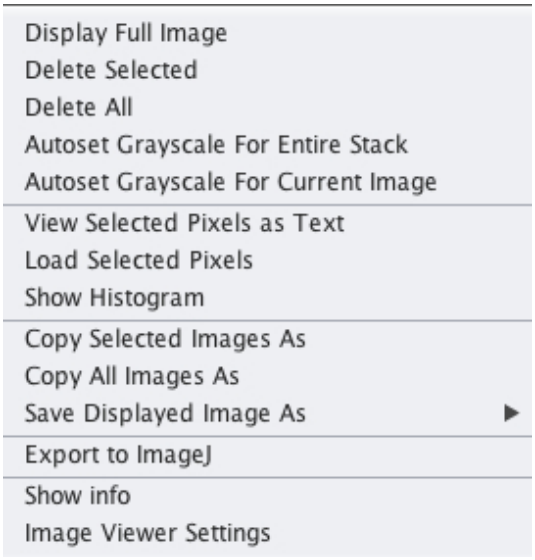

The right-mouse menu has many functions: images to be delete, the gray scale can be adjusted, pixel information can be displayed, Pixels can be viewed as text, look at a histogram. Finally, the image stack can be exported to ImageJ. This feature allows use all of the facilities in ImageJ to analyze and display images.

Figure 3.17: When the Image Viewer is running a right-mouse click brings up this menu. This menu can be used to delete images, load ROI pixels into the Ascii Data Viewer, export an image stack to ImageJ and a number of other useful functions.

Show Histogram will load the pixel values as a histogram.

Copy Selected Image As will copy a selected image stack to a new name and location.

Copy All Images As will copy all image stacks to a new directory.

Save Displayed Image As will save the currently displayed image as a jpg, png, tif or bmp image.

Export to ImageJ will export and open the image to in ImageJ.

Show info will display everything know about an image.

Image Viewer Settings will bring up a popup that configures the Image Viewer.

3.4.3.2 the Set Image area

The bottom-left area shown in Fig. 3.16 is used to set the image that is currently being displayed. This is done by simply moving either the slice or element number slider to the desired image. Alternately, the slice or element number can be entered and the viewer will display the desired image.

3.4.3.3 the Image Viewing area

The area in the center of the Image Viewer is used to display images. The Image Viewing area responds to a right-mouse click and the widgets on this submenu are the same as those shown in

Fig. 3.17 and we urge people to read the previous Subsection to determine their function. However, the image Viewing area has two menus that can be used to manipulate images. The first of these menus is at the top of the image viewing area. Here is a close up of this menu: Each of the buttons on

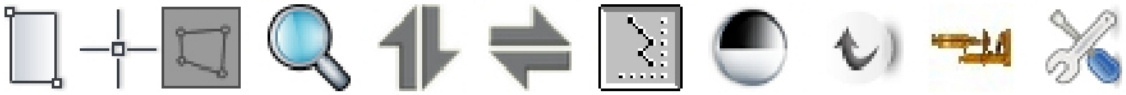

this menu performs various tasks associated with images. When activated, clicked on, the functions of these buttons are:

when activated a square ROI can be drawn on the image. Put the cursor in the image where an ROI is desired. Hold down the left-mouse and drag the cursor diagonally to draw the ROI. When an ROI is present the buttons on the right-hand side of the Image Viewer can be used to compute some statistics about values of the pixels. We will have more to say about this in Section 3.4.3.6

the star will create a point ROI. Simply place the cursor on the pixel to be capture and click the left-mouse button.

when activated the polygon widget draws a polygon ROI. To do this place the cursor where the ROI is to start and click the left-mouse button. A dot should appear at this vertex. Now move the cursor to the place where Move to the next polygon vertex to be and click the left-mouse button. Continue this for as long as needed. However, to close and end drawing the ROI end the drawing by placing the cursor on the starting vertex and left-mouse click a second time.

when activated this button will expand a square ROI. This button only functions when a square ROI is present.

when activated the image is flipped left to right.

when activated the image is flipped top to bottom.

when activated the pixel values contained within an ROI are averaged, written to an Ascii file and the Ascii file viewer is activated displaying the extracted pixels. This widget only functions when an ROI is present.

when activated the interface will attempt to determine a gray scale appropriate for the image stack. When the calculation is completed the images are displayed using the new gray scale.

when activated the original image is displayed using the original gray scale.

when activated the image stack is copied and exported to ImageJ.

when activated the user can set user preferences. This is the same popup that is available on the Settings/Preferences menu.

3.4.3.4 the Grayscale area on the bottom

The area at the bottom of the Image Viewing area is a grayscale area that can be used to manually set the grayscale. Here setting the grayscale means setting a lower and upper threshold on the pixel values that are displayed. To adjust the lower threshold, position the cursor on the grayscale bar at the value to be thresholded and do a left mouse click. This will move the red hatch mark to the cursor position, it will set the minimum threshold and finally it will redisplay the image using this threshold. Similarly, to set the upper threshold, position the cursor on the grayscale bar at the value threshold and do a right mouse click. This will move the blue hatch mark to the cursor position, it will set the maximum threshold and finally it will redisplay the image using this threshold. The threshold hatch marks, the blue and red vertical bars, can also be dragged. That is to say position the cursor on either the red or blue hatch mark and then drag the hatch mark to higher or lower values whichever is appropriate. Also the selected thresholds can be moved left or right on the grayscale bar. To move the threshold, located small tab in the center of the grayscale bar. This tab can be dragged left or right to vary the grayscale window that is displayed. This tab cannot be dragged until the left or right mouse clicks are used to raise or lower the displayed grayscale.

3.4.3.5 the Pixel Info area

The Pixel Information area is the area on the top right-hand side of the Image Viewer, see Fig. 3.16. It is used to display information about the image. The three widgets contained in this area have the following functions:

X Pos is the pixel number in the X (horizontal) direction. The first number is the pixel number in the raw image. Note that the graphic display in the Image Viewer is a 512 × 512 pixel display. The second number, the one in parentheses, is the pixel number in the graphics area and ranges from 0 to 511.

Y Pos is the pixel number in the Y (vertical) direction.

Value is the intensity of the pixel from the raw image at the current cursor position.

3.4.3.6 the Image Statistics area

is located on the right at the bottom of the Image Viewer. These widgets remain empty until one hits the "Get Statistics" button. When this button is activated the interface will compute a few basic statistics about the ROI if present, and about the entire image if no ROI is present. Here is a brief description of the information displayed in these widgets:

Mean contains the mean value of the images pixels in the selected region.

Max contains the maximum value of the image pixels in the selected region.

Min contains the minimum value of the image pixels in the selected region.

Sdev contains the standard deviation of the selected region. If p_i stands for the ith pixel in the selected ROI containing N pixels, then the "Sdev" value is calculated as:

$$\text{Sdev} = \sqrt{\frac{1}{N} \sum_{i \in \text{ROI}}^{N} [p_i - \text{Mean}]^2} \tag{3.1}$$

and Mean is the mean pixel value in the ROI and is given by

$$\text{Mean} = \frac{1}{N} \sum_{i \in \text{ROI}}^{N} p_i. \tag{3.2}$$

Note, for absolute value images this calculation should not be used as an estimate of the noise standard deviation. In a ROI containing only noise the Mean value is part of the noise and should not be subtracted when calculating the noise standard deviation. Additionally, even in an Absorption mode image this calculation may not give a good estimate of the noise standard deviation again because the deviation form the mean is not an estimate of the noise value.

RMS contains the **R**oot **M**ean **S**quare deviation of the selected region. Again if p_i stands for the ith pixel in the selected ROI containing N pixels, then the "RMS" value is calculated as:

$$\text{RMS} = \sqrt{\frac{1}{N} \sum_{i \in \text{ROI}}^{N} p_i^2}. \tag{3.3}$$

Note, for absolute value images this calculation should give a better estimate of the noise standard deviation and in absorption mode images this is the correct was to calculate the noise standard deviation.

Pixels contains the total number of pixels in the ROI.

Finally, there are three buttons at the bottom of the "Image Statics" area that are used to get and save various statistics computed by the interface:

Get Statistics will generate the statistics from the selected ROI. If no ROI is present then the statistics are generated for the entire image. Note this widget does not use the minimum or maximum pixel values set by the "Get Threshold Statistics" button, it simply computes the statistics for all pixels contained in the ROI.

Get Threshold Statistics will generate statistics for the selected ROI using a minimum and maximum pixel values. When this button is activated it will popup a widget in which a minimum and maximum pixel values can be entered. After entering these values selecting OK will cause the statistics to be calculated and updated. These updated statistics will ignore all pixels below the minimum or above the maximum. If no ROI is present then the statistics are generated for the entire image.

Save Statistics will bring up a popup that allows navigation and saving of the statistics. Here is an example of the saved statistics

```
Min             =       9.0002e+00
Mean            =       9.4447e+00
Max             =       9.9981e+00
SDev            =       2.8362e-01
RMS             =       2.8362e-01
Pixels          =       1292
```

and this output is left justified in the saved file.

Note all statistics generated are for a single image, there are no capabilities in the interface for applying the statistics calculations across multiple images.

3.4.4 Prior Viewer

The Prior Viewer is used by almost every package and, as its name implies, it is used to view, modify and generally set up the prior probabilities used in the Bayesian calculations. The viewer is shown in Fig. 3.18. In its appearance and function it is very similar to the Ascii Data Viewer, Fig. 3.12. Along the left-hand side is a list of priors that can be modified and on the right is a plotting area where the priors are plotted. In the example shown, it is the prior probabilities for the an inversion recovery model that are shown. Inversion recover models are single exponential plus a constant so there are three prior probabilities: the prior probability for the decay rate constant, and the initial and final intensities. When an item in list of parameters is activated with a mouse click, the prior probability for the selected parameter is displayed in the viewing area and the parameters that describe the prior probability are shown in the entry boxes just above the viewing area. The values in the entry boxes can be changed as needed. Selecting a different type of prior will cause the entry boxes to change to something appropriate for the selected prior. In the example shown, the prior is described by three parameters, a low, a peak and a high. The user can set these parameters to any values they wish and the interface will redisplay the prior as to reflect the changes. The interface will allows invalid values in these entry boxes, for example setting High less than Low, but it will not let the analysis to be run with invalid settings. If the parameters are invalid and the run button is activated, the interface will popup an error message and the error must be corrected before the analysis can be run.

When an analysis is run a list of the priors selected by the user is sent to the analysis package. The package normalizes the priors in such a was as to ensure the prior probability is always between zero and one and then proceeds to use the selected prior in the calculation. The normalization for the prior is set discretely, that is to say the prior probability is discretizing and the discrete prior is normalized so that the sum over the discrete samples is one. The prior is discretized on a 101 step inclusive interval. So for example, if the prior is an exponential of decay rate constant 1, with a range given by Low and High, then the normalization constant is given by

$$dX = \frac{[\text{High} - \text{Low}]}{100} \tag{3.4}$$

$$\text{Norm} = \sum_{i=1}^{101} \exp\left[-\left\{\text{Low} + dX(i-1)\right\}\right] \tag{3.5}$$

so the prior probability is given by

$$P(X|\text{High Low}) = \frac{\exp\left[-X\right]}{\text{Norm}}. \tag{3.6}$$

Plugging in zero for the low and 10 for the high, the normalization constant for the exponential prior is roughly 10.5. The prior starts at about 0.1 and goes down exponentially by about a factor of 22000.

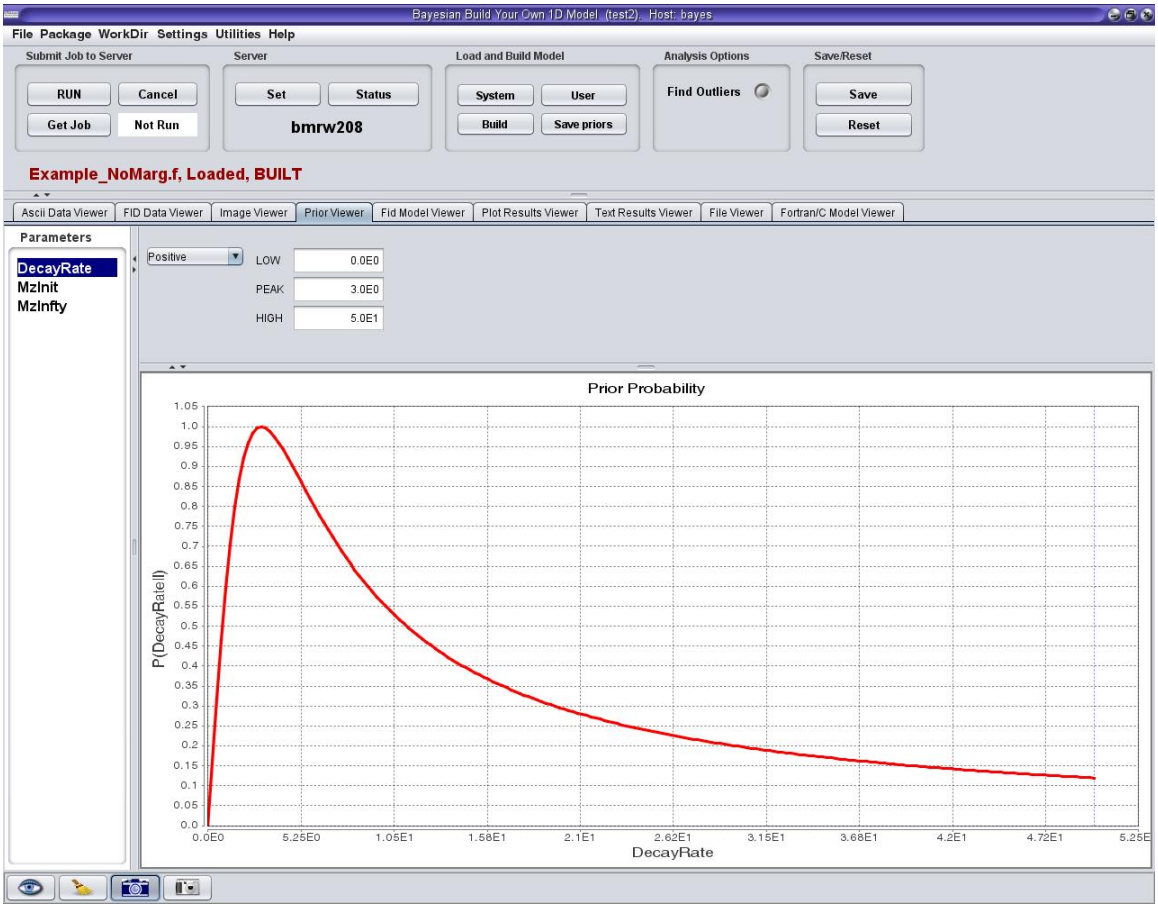

Figure 3.18: The Prior Viewer is used on most packages and it allows one to view and modify the prior probabilities used in the calculation. Here, the viewer is displaying a prior we call a positive prior. The positive prior has a number of very useful characteristics for scale parameters: It goes to zero at zero, thus preventing scale parameters from going negative. It is asymptotically Jeffreys' so allows the upper bound to be very large and it has a peak at a user specified location.

The prior selection box allows the user to select one of five different prior probabilities for a parameter. Ounce selected, the user can then specify the various elements that determine the prior shape. In all cases the user must supply a prior range, i.e., low and high parameter value. However, each prior type can have other characteristics that must be supplied by the user, for example a Gaussian requires a mean and standard deviation. Here is a description of the five priors and input requirements needed to generate a normalized prior:

Uniform selects a uniform prior probability having the user specified low and high parameter range. Uniform prior probabilities are not typically used for continuous parameters, although the user can certainly use them. Rather the interface typically uses a uniform prior probability for expressing an interference between models. For example, the Ascii Model Selection program can load up to 10 models. The prior probability for the model is uniform, i.e., one over the number of models.

Gaussian selects a Gaussian prior probability having the mean and standard deviations set by the user. Gaussian prior probabilities are the most common prior probability used in the software. Its a natural prior for amplitudes and it is typically used for any other parameter that can take on both positive and negative values.

Exponential selects a exponential prior probability having the decay rate constant set by the user. Exponential priors are typically used in the Bayesian Analysis software for discrete parameters. For example, the prior probability for the number of sinusoids in a model would typically be assigned a exponential prior.

Positive selects a positive prior probability having a peak value set by the user. An example of this prior is shown in Fig. 3.18. This prior is meant for scale parameters, and scale parameters can't be negative and this prior will not allow parameters to go to zero, let alone negative. This is clearly illustrated in Fig. 3.18, because the probability goes to zero as the value of the parameter goes to zero. This priors behavior is asymptotically Jeffreys' [32], again a characteristic desirable for scale parameters. This is also illustrated in Fig. 3.18, because for large parameter, the prior is dropping off very slowly; indeed like $1/X$. Finally, given that this prior goes to zero at zero and is asymptotically Jeffreys', the prior must have a peak value. This peak is also seen in Fig. 3.18, and as it turns out, the prior is essentially characterized by the location of this peak. The functional form of the prior is given by:

$$P(X|\text{Peak}) = \begin{cases} \left(\frac{1}{\text{Norm}}\right)\frac{X}{X^2+\text{Peak}^2} & \text{if Low} \le X \le \text{High} \\ 0 & \text{otherwise} \end{cases} \tag{3.7}$$

where X is the parameter, Low is the low parameter range, usually zero in this prior. High is the largest parameter value and Peak is the location of the peak. In the numerical calculations, Norm is set using the procedures described above.

Parameter is a prior used in the Bayesian Software package when a parameter is to be set to a constant. This prior is essentially a delta function and allows no variation in a parameter. As a computational note, the Markov chain Monte Carlo simulations do not vary any parameter that has a prior type set to parameter. There are a few packages that make use of this prior. For example, the enter Ascii Model package makes use of it as a means of passing

the calculation routines parameters. For example, a diffusion tensor model typically needs a conversion constant computed from some spectrometer settings. These spectrometer settings are passed to these modes as parameters and the models use them to compute the conversion factors.

As explained above, when these priors are used they are discretized and normalized to ensure that they sum to one. However, the Prior Viewer does not normalize these priors; rather the viewer sets the largest value in the prior to one. So care must be taken when comparing one prior to another because the prior scales shown in the viewer are not the same as those used in the calculation.

3.4.5 fid Model Viewer

Free induction decay data are time-domain data. However, most people look at fid data in the frequency domain. This presents a unique problem for frequency finding packages, because they work in the time-domain, but their outputs will almost certainly be viewed in the frequency domain. The frequency finding packages either generate time-domain models of the fid data directly or they generate a series of Ascii Model Files that can be used to generate a time-domain fid model of the data. Either way, the Fid Model Viewer must sometimes generate a time-domain fid model and then convert the time-domain fid model in the frequency domain for viewing. The Fid Model Viewer, shown in Fig. 3.19, is used to display the results from the analysis of an time-domain fid in the frequency domain. This viewer is activated when the Fid Model Viewer button is activated. If a time-domain model exists in the package, it is Fourier transformed, phased and the absorption mode spectrum is displayed. If no time-domain model exists, rather if Ascii Model Files exist, then the user must generate a time-domain fid model using the "Build BA Model For fid #" button. When activated the interface will dynamically generate a time-domain fid Model and then Fourier transform, phase and display the model in this viewer.

This viewer is very similar to the Fid Data Viewer discussed in Section 3.13. Indeed the underlying Java class that defines this viewer is the same as the Fid Data Viewer, here it has had a few additional functions added to it. In Section 3.13 we discussed the widgets along the bottom of the viewer and three of the widgets on the top part of the viewer, the "Trace" spinner, the "Data Type" selection and the "Options" widgets. The functions of these widgets is identical to their function on the Fid Data Viewer and, consequently, we refer the reader to Section 3.13 for those discussions. There are three main differences from fid Data Viewer: first the presence of the pull down menu containing the word "Trace". This widget allows different types of displays to be selected and we will give a description of the displays shortly. Second is the presence of the "Build BA Model For fid #" button. This button allows one to build and view a model fid and again we will have more to say about this shortly. Finally, the entry box just to the right of this button specifies the trace that is to be modeled.

3.4.5.1 The fid Model Format

Some packages generate fid models when they run: Big Peak/Little Peak and the Metabolite package both do this and some packages, Bayes Analyze and Find Resonances, generate Ascii Model Files that must be further processed to produce a fid model. The "Build BA Model" button builds these fid models from the Ascii Model Files when needed. The Fid Model Viewer button is present on all packages and can be used to generate Bayes Analyze models from previously analyzed spectroscopic

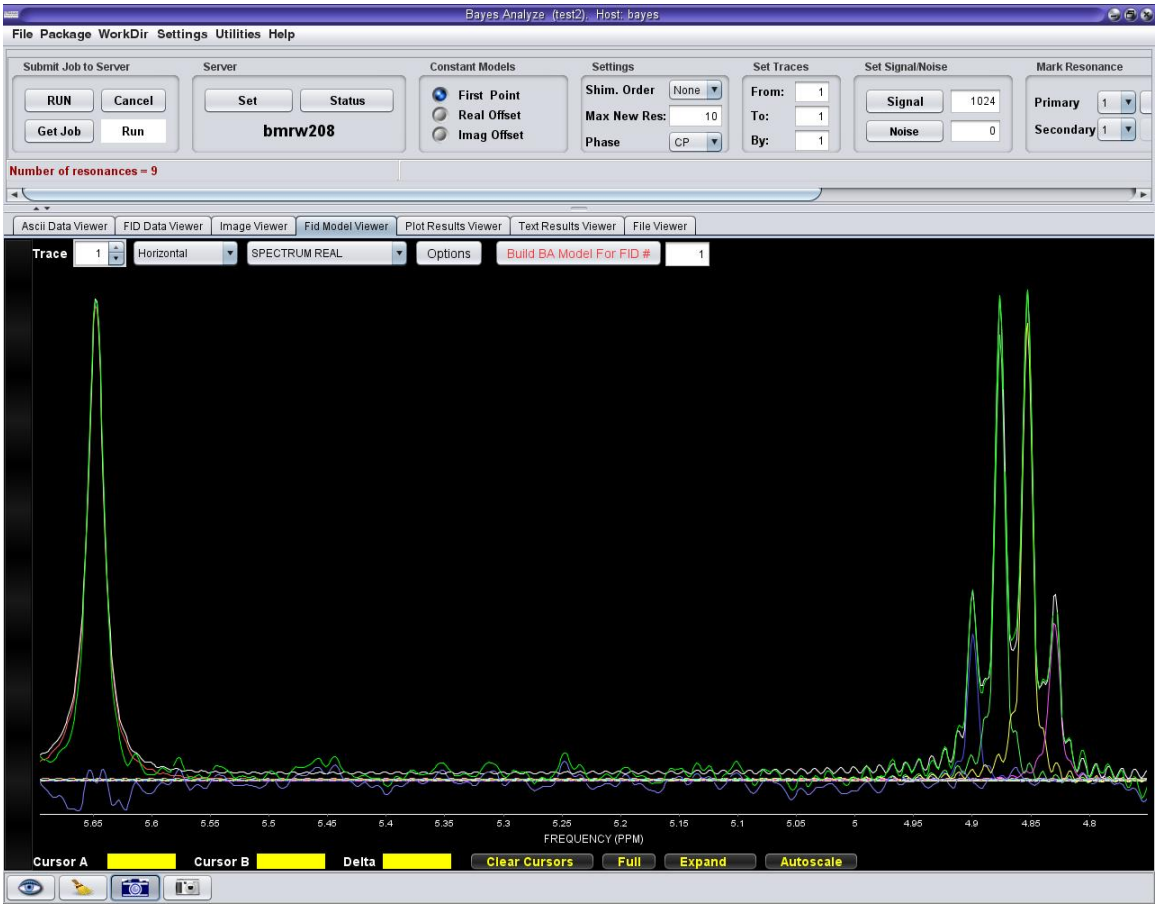

Figure 3.19: Frequency finding programs have the ability to display their outputs overlaid by the original fid. These outputs can be in the time domain or in the frequency domain. Additionally, because fid data can be arrayed displaying these outputs requires a special Viewer to generate and display the results. This is the job of the Fid Model Viewer shown here.

fid data. The interface uses the Ascii Model Files to generate a fid Model. Assuming Bayes Analyze Model files are present, when one clicks on the "Build BA Model fid #" button or one enters a fid trace number in the entry box then the interface sends a request to the current server to build a fid Model. The interface waits for this job to run. If needed the user is prompted for a password. When the job completes it is automatically retrieved by the interface and unpacked. The model fid is Fourier Transform and displayed in the Fid Model Viewer, see Fig. 3.19 for an example of the output. The Fid Model Viewer displays a number of traces, one each for each trace in the model field. The model fid is a time-domain arrayed fid contains the following traces:

Trace 1 is the complex time-domain fid data from the original input fid. i.e., the spectroscopic fid data that is being modeled.

Trace 2 is the complex time-domain fid model of the input fid.

Trace 3 is the complex time-domain difference between the model and the fid. This difference is usually called the residuals.

The remaining traces are the complex time-domain models of each individual resonance in the model. The resonances are in increasing frequency order, so trace 4 is the most negative frequency and the last resonance is the most positive. If the total number of resonances if N then the Model fid file contains $3 + N$ traces.

3.4.5.2 The Fid Model Reports

Next to the "Trace" spinner on the Model Fid Viewer, there is a widget that is initialized with the work "Trace" in it. This widget is the "Report Type" widget and the word "Trace" refers to the fact that the default report is to display the model traces. The report type widget is actually a selection menu and when activated it will display a selection of different kinds of reports that can be viewer/printed etc. This report type selection menu has the following selection items:

Trace is a display of the individual traces in the model fid. The trace spinner just to the left of the report type selection menu can be used to change the trace number. The number can also be changed by typing in the trace number to be viewed.

Data will display the original fid data. As noted the format of the fid data is also controlled by a selection menu. The default data type is to display the spectrum of the data.

Model will display the model fid data.

Residual will display the difference between the complex fid data and the model.

Vertical will display three plots stacked one above the other, the bottom plot is the original fid data, the middle plot is the model fid and the top plot is the residuals. Again the format of these displays is controlled by the data selection menu menu.

Overlay also displays the data in three stacked plots. Here the lower plot is the original fid data overlaid by the model. The middle plot is the residuals and the top plot is a plot of each resonance in the data.

Horizontal is the plot shown in Fig. 3.19. In this plot the data, model, residuals and the individual resonances are all plotted on the same scale. If examine Fig.3.19 is examined, it is almost impossible to distinguish the data and the model because they overlap each other almost perfectly. The residuals are the small randomly varying trace in purple. Note that these individual resonances do not actually reach the top of the spectrum in any of the peaks. None-the-less sum of these resonances fit the data perfectly.

Stacked is a plot of all traces in the model fid, one above the other. So the lower three traces are the data, model and residuals and traces 4 through the top are the individual resonances in increasing frequency order.

The type of data used in the displayed report is controlled by the data type widget. So for example by setting the data type to "fid" the report type widget will display all of its reports in the time-domain. Similarly, if the data type is set to "Spectrum Real" then the reports use the real part of the spectrum of the model.

3.4.6 Plot Results Viewer

The Plot Results Viewer is for all practical purposes the Ascii Data Viewer described in Section 3.12 and everything said about the Ascii Data Viewer applies here. In this Subsection we are going to discuss the differences between the Ascii Data Viewer, explain briefly how to use this viewer and give a general introductions to the kinds of plots this viewer displays. In particular we are going to explaining a few of the more common types plots that show up in this viewer. However, each package may have plots unique to that package and if a plot is discovered that is not explained here, then consult the Chapter on the package being used.

The Plot Results Viewer is shown in Fig. 3.20 and if Ascii Data Viewer is compared to the Plot Results Viewer, one will find that about the only difference is the label on top of the plot files list. In the Ascii Data Viewer this label reads "Ascii Data" while here it reads "Output Plots". In the Ascii Data Viewer the area under this label contains a listing of all of the files that have been loaded into the experiment, while here it contains a listing of all of the output plots generated by the package. The example shown in Fig. 3.20 is the result of running a biexponential model with a constant on biexponential data. In the output plots list there are 15 different plots. When a plot is activated, the output plot list entry is highlighted and the plot is displayed. The plots are organized roughly as three groups of plots, in the top part of the output plot list are plots showing how the model fits the data. In the middle section are the posterior probabilities for the various parameters appearing in the model. Finally the bottom contains a number of plots that are meant as aids in understanding the outputs from the simulations.

Like the Ascii Data Viewer, this viewer also responds to a right mouse click. This is true for the plotting area and the Output Plots area. However, the menu shown in the plotting area is the same as in the Ascii Data Viewer, but the menu shown in the Output Plots differs. Here is a brief description of this submenu and what it does:

Plot Information will display all of the information available to the interface about this plot. Figure 3.21 is an example of what is displayed when the plot information menu item is selected. We have numbered the lines in this figure to make referencing them easier. Line 01 indicates that this is a line plot and we are to plot column 1 vs 2 from the file named in Line 03. Line 02 is an internal indicator and it tell us that column 1 of the Bayes.Mcmc.Samples file is to be

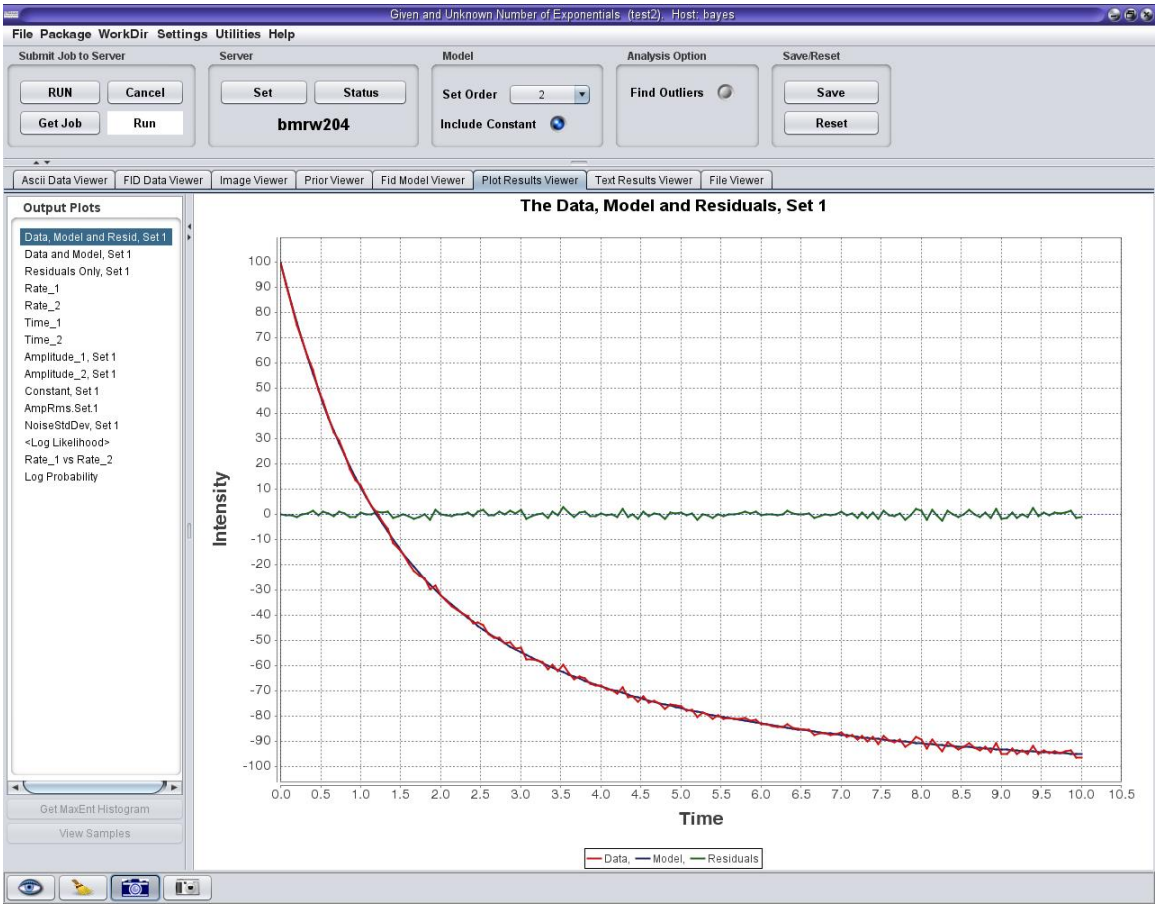

Figure 3.20: Ascii packages and some fid packages output Ascii Plot files/ These plots are displayed by the Plot Results Viewer. Roughly speaking there are six different types of these plots: plots involving the data, model and residuals; plots involving the posterior probabilities for parameters; plots showing correlations between parameters and two types of plots that are used to determine if the analysis ran/converged correctly.

```
01  Line 1 2
02  1
03  BayesExpGiven.Rate_1
04  Prob for the Rate_1, Given a 2 Exp Model and a Const
05  Rate_1
06  Rate_1
07  Probability Density
08  Mean: 0.26448; Sd: 2.15E-02; Peak: 0.26487
```

Figure 3.21: The plot information widget will popup up a window containing information about the currently displayed plot. We have numbered the lines in this display to make referencing them easier. Line 01 indicates that this is a line plot and that column 1 vs 2 from the file named in Line 03 is to be plotted. Line 02 is an internal indicator and it tell us that column 1 of the Bayes.Mcmc.Samples file is to be displayed should the user activate the "View Samples" button. Line 03 is the name of the file being plotted. The file is located in the "Bayes Home director/WorkDir/BayesOtherAnalysis" directory. Line 04 is the title of the plot. Line 05 is a long abscissa label. Line 06 is a short abscissa label. Line 07 is the Y axis label. Finally, Line 08 on probability plots, contains the parameter estimates.

displayed should the user activate the "View Samples" button. Line 03 is the name of the file being plotted. The file is located in the "Bayes Home director/WorkDir/BayesOtherAnalysis" directory. Line 04 is the title of the plot. Line 05 is a long abscissa label. Line 06 is a short abscissa label. Line 07 is the Y axis label. Finally, Line 08 on probability plots, contains the parameter estimates information displayed as part of the title.

Show Data From Plot will display the data actually plotted in a popup. This display is usually two column, but it will be multicolumn when plots having more than a single trace are displayed.

Show Source File for Plot will popup a window containing the original source file. This source file can be different from that displayed by the "Show Data From Plot" widget. For example, the file containing the data, model and residuals is a four column file; while a plot of only the data and the model would contain only three columns even though the source file is four columns.

3.4.6.1 the Data, Model and Residuals Plots

It was mentioned early that the output plots list is roughly divided into three groups of plots. The first group of plots, plot the the data, model and residuals in a number of different ways. This group plot generally consist of three different plots: one of the Data, the model and Residuals; one plot of the data and the model and finally a plot of just the residuals. Figure 3.20 is the result of running a biexponential model with a constant on biexponential data. The plot shown in this figure was generated by clicking on the "Data, Model and Resid, Set 1" line in the output plot list. This plot contains three traces, the data in red, the model in blue and the residuals in green. These three plots are generated from the Markov chain Monte Carlo simulation that had maximum posterior

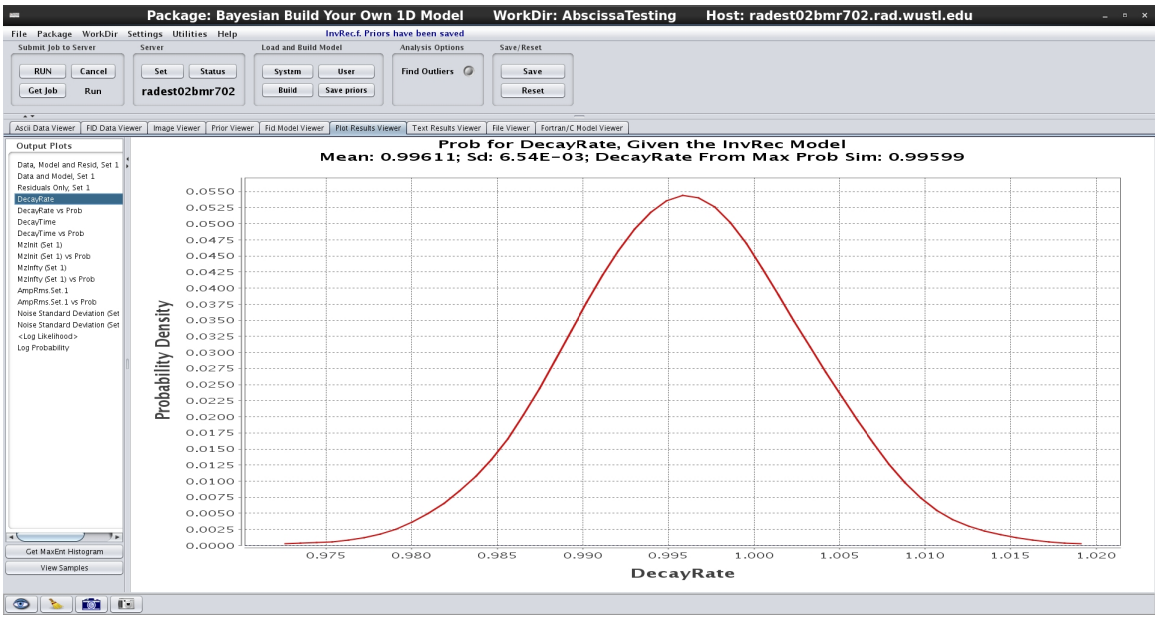

Figure 3.22: The whole purpose of the Markov chain Monte Carlo simulation is to derive an approximation of the Bayesian posterior probability for the parameters appearing in a model. To do this, a Markov chains is used to gathers samples from the joint posterior probability for all of the parameters, and then uses Monte Carlo integration to obtain samples from the posterior probabilities for the individual parameters. These samples are then smoothed, binned and displayed in plot like the one shown here. This particular plot is for the posterior probability for the smallest decay rate constant in a biexponential model containing a constant.

probability. If multiple data sets were processed then the output plot list area will contain multiples of these data, model and residual plots, one set for each data set.

3.4.6.2 the Posterior Probabilities Plots

All of the Bayesian analysis packages, with two exceptions, use Markov chain Monte Carlo to approximate the Bayesian posterior probability for the parameters appearing in the model. To do this, a Markov chain Monte Carlo simulation is used to draw samples from the joint posterior probability for all of the parameters appearing in the model and Monte Carlo integration is used to obtain samples from the posterior probabilities for the individual parameters. We display these samples in several different ways: This default view of these samples is as a smoothed binned histogram like the one shown in Fig. 3.22. This particular plot is for the posterior probability for the smallest decay rate constant in a biexponential model containing a constant. However, the samples themselves can be viewed by activating the "View Samples" button at the bottom of the output plot list area.

The default binned histograms are generated vary crudely and are only meant as an aid in understanding things like the mean, standard deviation, and symmetries of the samples, they are not used for any calculations. Indeed, in the output reports all parameter estimates, means, peaks

and standard deviations are computed directly from the samples. To generate these smoothed binned histograms, a Gaussian is placed in the histogram at the point where a sample occurred. This Gaussian has a standard deviation such that it goes through one e-folding every two bins in the histogram. The Gaussian is then evaluated over a symmetric ± 6 bin interval, so 13 bins are included, and the value of the Gaussian is added to the histogram. Finally, the histogram is normalized so that the sum over all bins is one. This rather crude histogram technique works well provided the histograms are reasonably smooth.

One can replace a smoothed binned histogram by a fully Bayesian estimation of that histogram. This MaxEnt histogram is obtained by activating the "Get MaxEnt Histogram" button near the bottom of the output plot list. Activating this button will send a request to the selected server to run the MaxEnt Histogram package. The interface runs the request in background so continue working on the interface. If the same smoothed histogram is being displayed when the MaxEnt histogram package finishes, the smoothed binned histogram is replaced by the MaxEnt histogram and the histogram is redisplayed. If a different smoothed histogram is displayed, the MaxEnt histogram overwrites the appropriate smoothed histogram and the MaxEnt histogram will be displayed the next time the user displays the histogram.

The MaxEnt calculation is as computationally intensive as running many other package, so only generate a MaxEnt histogram when a better characterization of the samples is really needed. Because this calculation is a full Bayesian calculation, the resulting histograms have error bars. For a detailed description of these Bayesian calculations see Chapter **??**.

3.4.7 the Posterior Probability Vs Parameter Samples plot

In release 4.10 a new type of plot was added, see Fig. (3.23). For each plot of the posterior probability for a parameter, there is an addition scatter plot of the parameter verses the posterior probability. If the posterior probability for a given parameter is highly symmetric having a well defined maximum, then this new plot, Fig. (3.23), will look almost exactly like the plot shown in Fig. (3.22). In this new plot the horizontal axis is the value of the parameter and the vertical axis is the unnormalized posterior probability. When the Markov chain Monte Carlo simulations are running it is the logarithm of the posterior probability that is computed. To obtain the unnormalized posterior probability shown in this plot, we locate the sample that has maximum posterior probability, subtract that value from all of the other samples and then exponentate. Finally, the samples are normalized so that the peak sample is the same as the peak shown in Figure 3.22. When the posterior probability is well peaked, these plot is almost identical to the plot shown in Fig. (3.22).

One of the reasons this plot was added was to help eliminate a common confusion. In the heading of the plots of the posterior probability, part of that heading reads ""Parameter Value" From Max Prob Sim: xxxx"" where "Parameter Value" is the name of the parameter being shown, for example in the exponential package it would read "Rate_1" and "xxxx" is the value of the parameter for which the posterior probability was maximum, in the example shown this value is 0.99599. For plots like the one shown, the peak in the histogram and the parameter that maximized the posterior probability are essentially the same. However, it is possible for the parameter value at the peak of the histogram, and the parameters from the simulation having maximum posterior probability to be very different.

For example, in Fig. 3.24 the resulting histogram and the samples have been plotted on the same scale. The histogram is shown as the line in this plot and the samples are the open circles. In this example, the peak of the histogram is about 0.1, while the sample that had maximum posterior

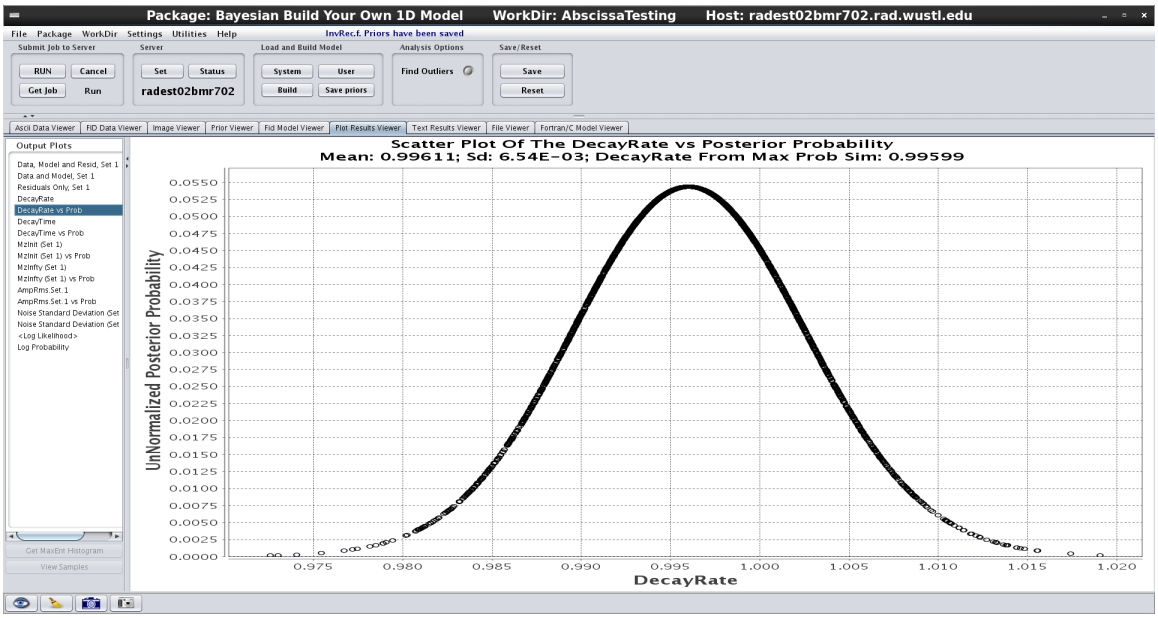

Figure 3.23: In addition to plotting the posterior probability for the parameter the unnormalized posterior probability for the parameter is plotted against the parameter value used in evaluating the posterior probability. As illustrated here, for well peaked posterior probabilities, these samples will look almost exactly like the binned histogram, Fig. (3.22).

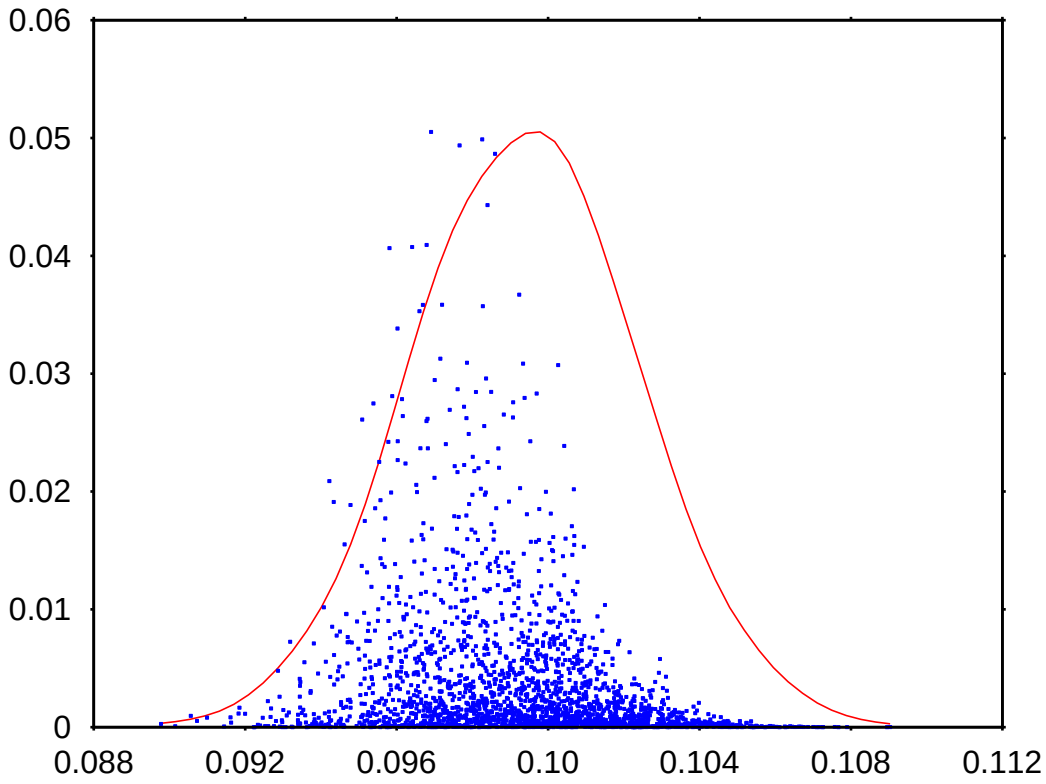

Figure 3.24: When the peak of the posterior probability is not symmetric or only weakly a function of a given parameter, then it is possible for the parameters of the simulation that had maximum posterior probability to be significantly different from the mean of the Markov chain Monte Carlo samples. This is illustrated here, the peak of the histogram is rather different from the sample that had maximum posterior probability.

probability is somewhere around 0.094. So be warned the peak in the histogram and the sample that had maximum posterior probability can be very different.

3.4.7.1 the Expected Log Likelihood Plot

A plot of the expected logarithm of the likelihood as a function of the annealing parameters is shown in Fig. 3.25. When the Markov chain Monte Carlo simulations start, the annealing parameter, labeled Beta, starts at zero and is increased to one according to some annealing schedule. The Markov chain Monte Carlo simulations are being done using simulated annealing. In simulated annealing one raises the likelihood to a power β. If the parameters of interest are designated by M, and the data as D, then logarithm of the posterior probability is given by

$$\log P(M|DI) = \log P(M|I) + \beta \log P(D|MI) \tag{3.8}$$

where $P(M|DI)$ is the posterior probability for the parameters given the data and the prior information, $P(M|I)$ is the prior probability for the parameters, β is the annealing parameter and $P(D|MI)$ is the direct probability for the data given the parameters M and the prior information I and in this discussion it is called a likelihood. When the annealing parameter is zero, only the logarithm of the prior probability contributes to the calculation and the simulations explore the prior probability. Because the data are not being used, the fit to the data is very bad and, consequently, the expected logarithm of the likelihood starts out at a very low number. As the annealing parameter is increased, the logarithm of the likelihood contributes more and more to the calculation and so the expected logarithm of the likelihood increases coming to a maximum when the model has fit the data as well as possible. All the packages that run Markov chain Monte Carlo simulations run multiple simulations in parallel. The expected logarithm of the likelihood shown in Fig. 3.25 is the average of logarithm of the likelihood over all of these simulations. See Subsection B for a more detailed discussion of simulated annealing and Markov chain Monte Carlo.

3.4.7.2 the Scatter Plots

There are two additional plot types that need to be discussed: scatter plots and the logarithm of the posterior probability plot. Scatter plots are generally toward the bottom of the output plot list and have have names of the form "parameter 1 vs parameter 2", where parameter 1&2 are any two parameters from the simulations. An example scatter plot is shown in Fig. 3.26. This particular example was generated using the exponential package with a biexponential model plus a constant. Many samples are drawn from the joint posterior probability for the parameters. In the example shown here, there were 50 simulations and 30 samples were drawn from each simulation, so there are a total of 1500 samples. The scatter plot shown in Fig. 3.26 is of Rate_1 vs Rate_2. Each point in this plot is one of the 1500 simulations. The coordinates of each point are the values of the decay rate constants in a given simulation.

Scatter plots can look like an ellipsoid that is aligned with the axes, in which case the parameters are uncorrelated. Scatter plots can also be tilted, like the one shown in Fig. 3.26 in which case the parameters are correlated. Finally, scatter plots can take on highly irregular shapes and can even have sharp cutoffs when parameters have natural bounds. The scatter plot shown in Fig. 3.26 is an almost classical example of a scatter plot of two correlated parameters. Note the almost linear like upward scatter in this plot. This type of feature is indicative of a correlation between the sum and difference of the decay rate constants. The density of points in a scatter plot is a sample from the

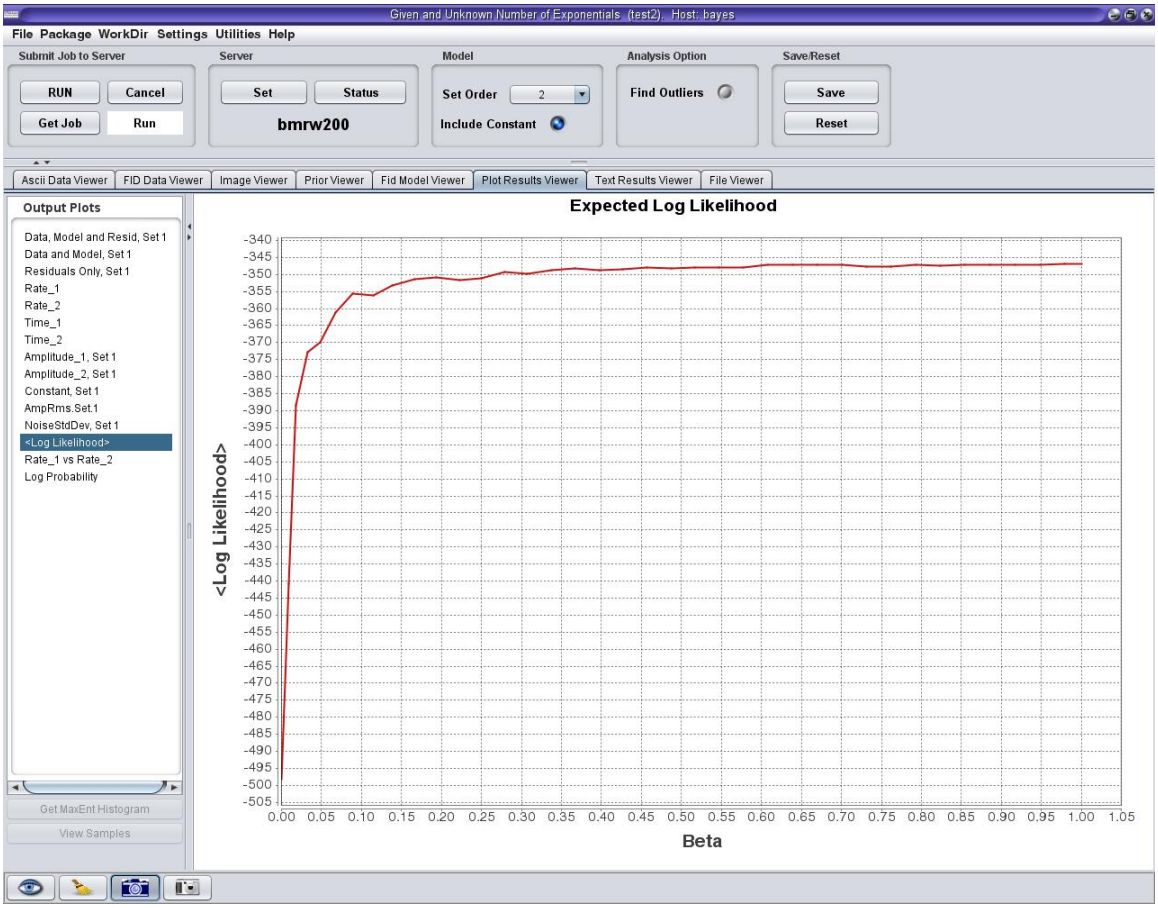

Figure 3.25: When the Markov chain Monte Carlo simulations run the intermediate results can be used to compute the logarithm of the posterior probability. Additionally, a plot of the expected logarithm of the likelihood can be produced. This plot can be used to aid one in determining if the Markov chain Monte Carlo simulations have converged.

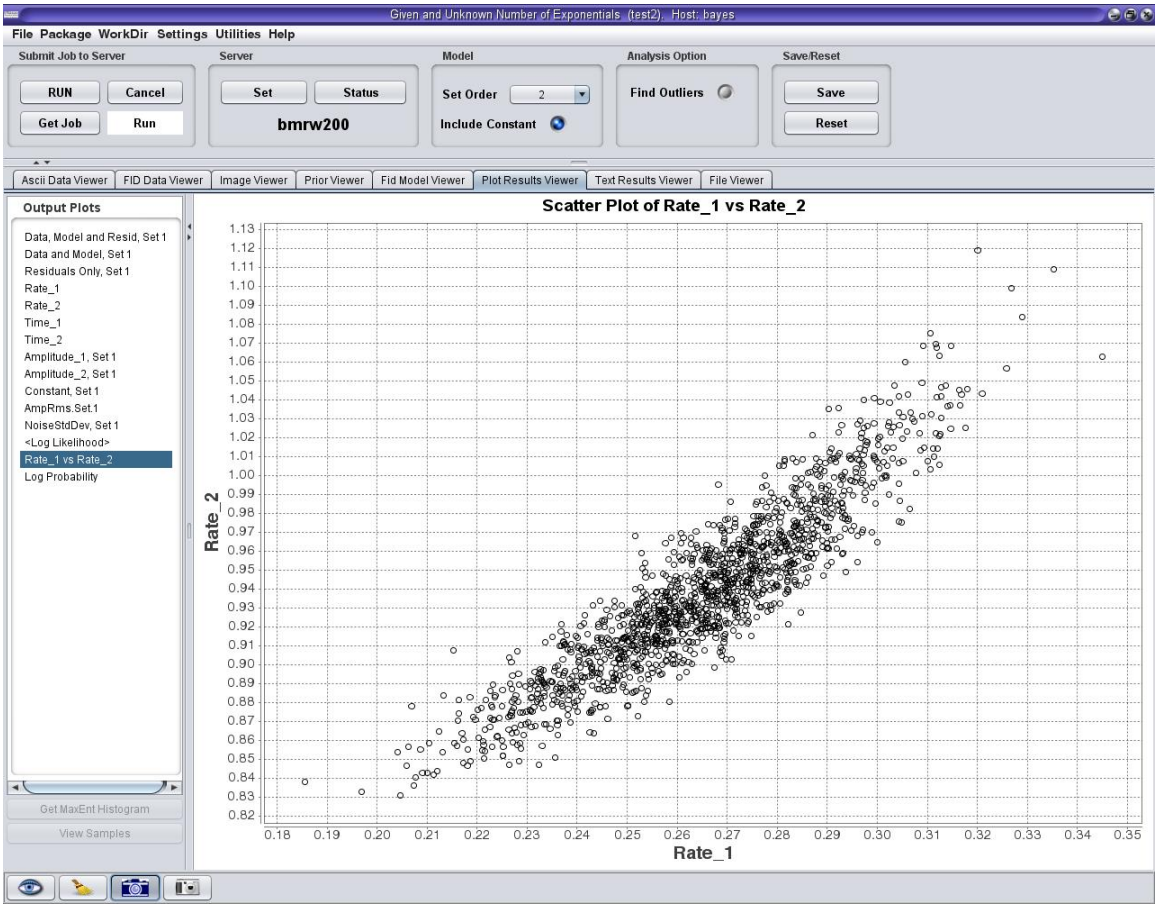

Figure 3.26: When the Markov chain Monte Carlo simulations run, many of the packages output scatter plots that are used to determine if the parameters are correlated and to check on the convergence of the simulations. Here is an example scatter plot created using the exponential package with biexponential data with a constant. The plot is of the two decay rate constants one verses the other. In this case there is a strong correlation between decay rate constants, but the simulations are otherwise well converged.

joint posterior probability for the decay rate constants. In Fig. 3.26 the scatter plot is elongated along the vector sum of the decay rate constants and it is contracted by about a factor of 3 along the vector difference. This implies that the sum of the decay rate constants is not well determined compared to the difference. This type of correlation is very common among models containing multiple exponentials.

Not all packages produce scatter plots and even packages that do generate scatter plots do not always generate scatter plots of all of the parameters. For example here, there are three amplitudes and two decay rate constants in the model, but only a scatter plot of the decay rate constants was output. That's because the exponential package uses a marginal posterior probability and the amplitudes are not varied by the Markov chain Monte Carlo simulations. Scatter plots are only produced for parameters that are varied by the simulations. Additionally, some packages have a very large number of parameters and because the number of scatter plots increases like the square of the number of parameters, scatter plots are typically not output when the number of these plots would be very large.

3.4.7.3 the Log Probability Plot

The last plot that we are going to discuss is shown in Fig. 3.27. This plot is a bit of a mess, and generally speaking, the messier it is, the better. In the Markov chain Monte Carlo simulations used by the Bayesian Analysis software multiple chains are run in parallel. In the annealing phase, the simulations are not in equilibrium. However, when one reaches the sample gathering phase, the simulations are should be in equilibrium. Here, equilibrium means that as the Markov chain Monte Carlo simulations are run the details of each simulation change, but the expected values of the parameters, probabilities, likelihoods etc. remain constant, i.e., the system is at a static equilibrium. In the literature this condition is often called detailed balance.

In the sample gathering phase, the simulations are run through a predetermined number of steps and then each simulation is save. This process is repeated until all of the user specified repeats have been gathered. Each line in the plot shown in Fig. 3.27 is the logarithm of the posterior probability for one of the simulations as a function of repeat number. As the simulations are run, the posterior probability for a given simulation is increasing and decreasing as a function of the repeat number. When all of the simulations are doing this, the simulations are often said to be mixing. Indeed in this plot, the simulations are mixing so well that it is impossible to follow the trajectory of a single simulation. However, note there is a fairly sharp maximum that bounds the posterior probability. This bound reflects the fact that the model can only fit the data so well. Individual simulations come up to the maximum and then move away from it. Here, the size of the deviations from this maximum are on the order of a few e-folding, with the maximum deviation being about 10 and a typical deviation being more like one or two. The size of the deviations and the presence of a sharp boundary are both data and model dependent. However, a good rule of thumb concerning the size of these deviations is that each parameter in the model can, on average, cause a deviation of one and perhaps two e-folding and if there are 2 parameters then an average deviation would be 2 to 4 e-folding. Models with more parameters will deviate from the maximum by roughly the number of parameters times one or two. Large deviations caused by a single parameter are hard to obtain simply because such a large deviation in the posterior probability is rejected by the Markov chain.

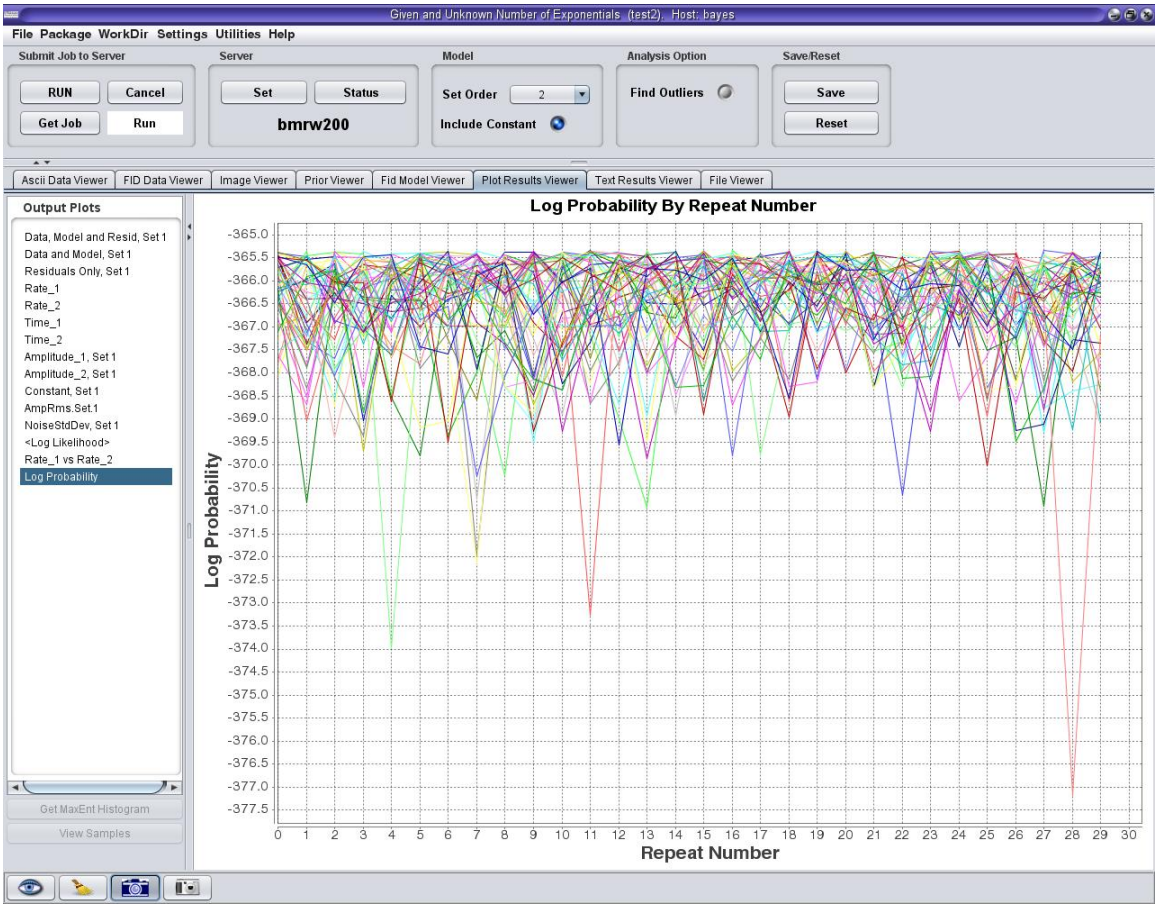

Figure 3.27: When the simulated annealing phase is finished for all of the simulations, the program that run the simulations begin gathering samples of each of these simulations. The program tries to gather uncorrelated, i.e., independent samples. If the user specified 50 simulations and 30 samples, the the program tries to gather 30 independent samples of the 50 simulations. The shown here is the logarithm of the posterior probability for each simulating plotted by sample number.

3.4.8 Text Results Viewer

After an analysis has been run the viewers are used to look at the outputs from the analysis. The Text Results view show in Fig. 3.28 is a typical example of the Ascii outputs. We are going to describe each of these shortly but note that the Text Results Viewer shows the output from two different packages: The current package, labeled standard, and Bayes Analyze. The reason the Bayes Analyze package is shown is because that package is unique in that some of the outputs from that package can server as inputs to the current Ascii package. Consequently, the current working directory can contain outputs from the current package in BayesOtherAnalysis and outputs from Bayes Analyze in BayesAnalyzeFiles. Activating any of the items on these selection menus will cause the appropriate file from the selected package to be listed in the Viewing area. If the requested file does not exists then an appropriate message is displayed. You can also switch reports by selecting them or by using the up and down arrows on your keyboard. However, you cannot jump between the standard output and the Bayes Analyze outputs using the up and down arrows. The entries in both text results viewers, standard and Bayes Analyze, are fixed and do not vary from package to package. Along the top of the viewing area are a number of buttons that are specific to the Text Results Viewer. Here is a description of these buttons:

Print will direct the currently selected file to the printer. the button will popup a widget that allows you to select a few print options including the printer.

Copy will copy the currently selected file to your clipboard.

Save will save the current copy of the file on top of the original file located in BayesHome/CurrentWorkDir/BayesAnalyzeFiles.

Save As will popup a navigation window that will allow you to navigate to the location you wish the file to be saved and then it will save the file.

Enable Editing will allow you to change the contents of the viewing window. For example, you might want to specialize a title or some text in the window as a reminder of what you did in the analysis prior to saving the results. When activated the check box just to the right will toggle on and off. You can also just check the box to enable editing.

Scroll Up will cause files to be positioned at either the beginning or ending of the file.

Settings will popup a window that allows you to configure the Text Results Viewer. Mostly it allows you to set the fonts and font point sizes.

The Standard selection list has a number of widgets associated with it and these reports are briefly described here. See the individual packages for a more detailed description of these reports. In all cases when a widget in the standard selection list displays a file it is located in your BayesHome/CurrentWorkDir/BayesOtherAnalysis directory.

Instructions will redisplay the instructions that are shown whenever a package is started. These instructions are typically terse, but they will indicate the types of things that must be done to successfully run a package. For more information on the individual packages consult the appropriate Chapter.

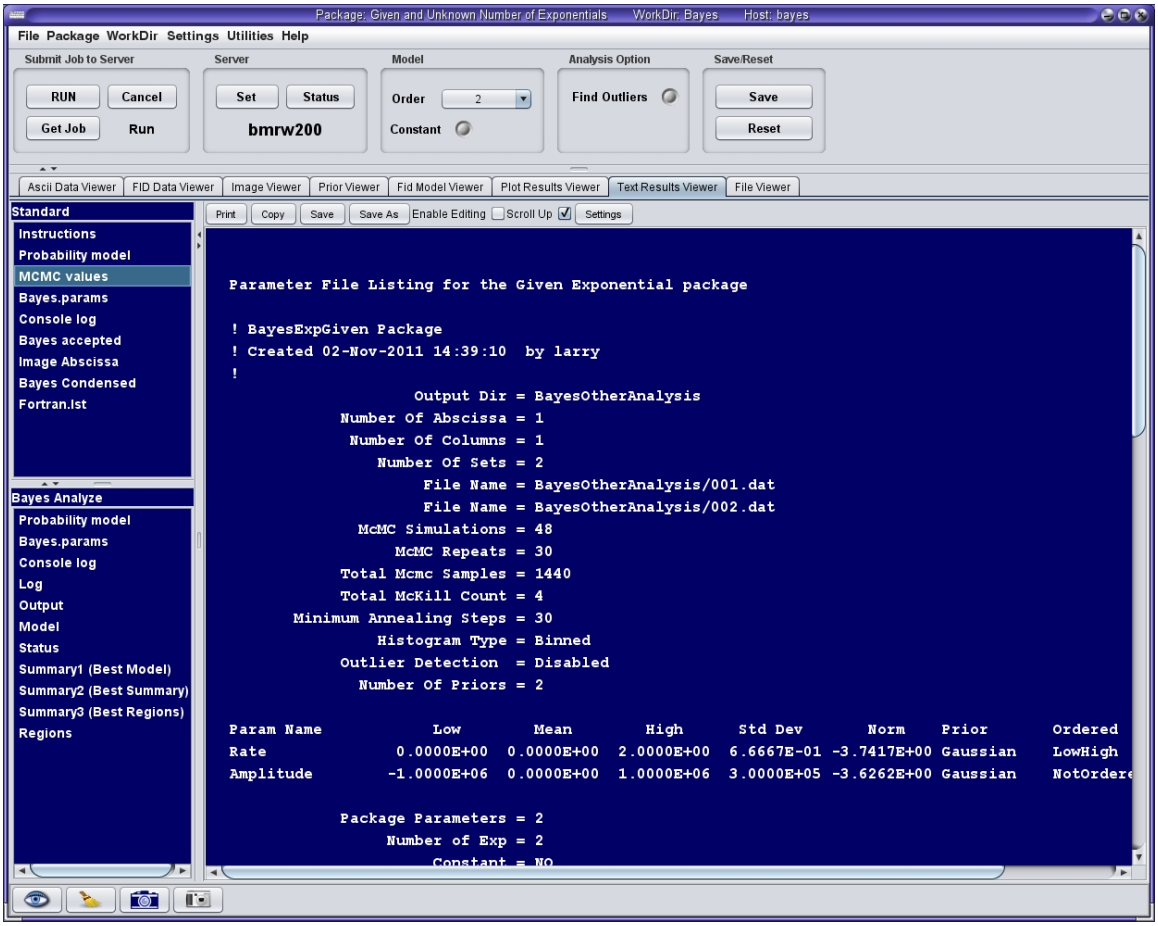

Figure 3.28: After a package has been run, the various outputs can be viewed using the Viewers. The Text Results Viewer is used to view the various reports that are output from the packages. The area along the left-hand side lists the various reports available from the current package, top, and from the last run of Bayes Analyze, bottom. The text has a brief description of the output and you must consult the package Chapter to find out about package specific reports.

Probability model will display the contents of the Bayes.prob.model file. This file usually contains one record for each time an analysis package has been run and it contains the results of a thermodynamic calculation for the probability for the model. See [55, 36] for more on thermodynamic integration.

Mcmc Values will display the contents of the PackageName.mcmc.values file, where "PackageName" is our internal name of a package. If you are looking for this file, there is only a single file in the BayesOtherAnalysis directory having the suffix "mcmc.values". For a much more detailed description of this report, see Subsection D

Bayes.params will display the contents of the Bayes.params file. The Bayes.params file contains a complete description of the package setup including the prior settings.

Console.log will contain all outputs that went to console when a job is running. Usually this is not useful, but when something goes wrong the console log would be consulted.

Bayes.accepted is the file that is displayed when the "Get Job" widget is activated and the job has not yet finished. The exact contents of this report are specific to a package but in general terms it consists of two parts: A header:

```
Accepted Report given the ExpTwoConst_Marg model

                 McMC Phase:       Sampling
          McMC Simulations:             50
              McMC Repeats:             30
   Number Killed Per Cycle:              5
       Min Annealing Steps:             51
              Current Step:             30
 Fraction Samples Gathered:        1.00000
Average Log Posterior Prob:-3.66274138E+02
 StdDev Log Posterior Prob:        1.05944
         Average Log Prior:        -19.245
    Average Log Likelihood:       -347.029
     StdDev Log Likelihood:          1.055
```

which is reasonably standardized. In general terms most of this header is either setup information concerning the package, or its a set of current statistics about the current status of the package. For example the Simulations, Repeats and Minimum Steps are all setup parameters; while the others are current status information. For example the various average probabilities are the mean value of the given probability averaged across the number of simulations, in this case 50. Similarly, the StdDev values are the standard deviation of the given probability. In general terms the likelihood should increase as a function of the annealing step; while the posterior probability will decrease. In both case the standard deviations of these quantities will be large for low values of the annealing parameters and will become smaller as the annealing parameter increases.

The lower part of the Bayes.accepted file is package specific and the number of entries in this part of the report can be highly variable. In general terms the lower part of the report will

contain some statistics about how often each parameter in the Markov chain Monte Carlo simulation is being accepted and rejected. Here is an example taken from the Enter Ascii Model package with the marginalized two exponential plus a constant model

```
Param Desc       Avg. Param.    Param Sd.    Proposal    Prior Contrib     Rate
DecayRate1     2.682716E-01  2.179858E-02  4.367748E-02 -4.525808E+00    0.2393
DecayRate2     9.382104E-01  4.349060E-02  7.788657E-02 -3.835734E+00    0.2485
```

Because this is a marginal probability density function only the decay rates are varied in the McMC simulations and consequently only the decay rates are shown in the accepted report. The entries are the average parameter value, its current standard deviation, the current proposal used in the McMC simulation, the contribution of this parameter to the logarithm of the prior probability and finally the acceptance rate. The acceptance rate, labeled Rate, is the number of times a parameter was accepted divided by the total number of times a new value was proposed. The packages try to keep this acceptance rate between 20 and 30%. For more on how this is done see Section B and for more about the Bayes.accepted report generated by a specific package consult the appropriate Chapter.

Image Abscissa will display the file BayesHome/WorkDir/images/Abscissa if its available. This file is used in image processing to tell the Ascii model files what the abscissa values are for the Ascii model. Consult the Section A for more about this file.

Bayes Condensed displays the file BayesHome/WorkDir/BayesOtherAnalysis/Bayes.Condensed.File an example of this file is shown in Fig. 3.29. The file consists of one line per parameter in the model. Each line contains the parameter name, the mean and standard deviation of the parameter computed from the simulations and the parameter value taken from the simulation that had maximum posterior probability. The header shown in this figure is not present in the actual condensed file, that header was put in the figure only as an aid in identifying the various fields.

The names shown in this plot are generally the names assigned to the parameter by the user or by us when we wrote the code. These names should be simple and self explanatory. As illustrated in Fig. 3.29, the interface will modify the names of the amplitudes and noise standard deviations by appending a data set number to them when multiple data sets are used. In this example the data set numbers are the ".01" and ".02" suffixes. So, for example, "BayesExpGiven.Amplitude_2.02" is the amplitude assigned to decay rate 2 in data set number "02". Similarly, "BayesExpGiven.Amplitude_2.01" is the amplitude number 2 in data set "01".

Fortran.lst will display the file BayesHome/WorkDir/model.compile/CurrentModel.lst where "CurrentModel" is the name of the Fortran of C model loaded. Note that this listing will contain any errors issued by the Fortran or C compilers when the "Build" button is activated.

Note that if you were to use the File Viewer to look at the contents of a WorkDir after it has been successfully run, you will find that it contains many more files than mentioned here. Those other files contain the probability density functions and the other reports mentioned earlier in this Section.

As noted, the Text Viewing area allows you to view the output from the current model, but it also lets you view the Bayes Analyze output from the previous run of Bayes Analyze. This selection menu also contains a fixed number of entries. Activating each entry will do the following:

Parameter Name	Mean	StdDev	Peak
BayesExpGiven.Rate_1	1.63255E-03	1.47866E-03	2.22598E-05
BayesExpGiven.Rate_2	2.95738E-01	6.03008E-03	3.01355E-01
BayesExpGiven.Time_1	4.44296E+03	5.02856E+04	4.49241E+04
BayesExpGiven.Time_2	3.38280E+00	7.03289E-02	3.31834E+00
BayesExpGiven.Amplitude_1.01	-1.01643E+02	1.84909E+00	-9.97170E+01
BayesExpGiven.Amplitude_2.01	1.01287E+02	1.72691E+00	9.94602E+01
BayesExpGiven.AmpRms.Set.1	1.43493E+02	2.52340E+00	1.40840E+02
BayesExpGiven.NoiseStdDev.01	1.08257E+00	8.47944E-03	1.07244E+00
BayesExpGiven.Amplitude_1.02	-1.00316E+01	1.80558E-01	-9.84339E+00
BayesExpGiven.Amplitude_2.02	9.70180E+00	1.69333E-01	9.52261E+00
BayesExpGiven.AmpRms.Set.2	1.39556E+01	2.46988E-01	1.36957E+01
BayesExpGiven.NoiseStdDev.02	1.07327E+00	5.41963E-04	1.07274E+00
BayesExpGiven.RmsAmplitude_1	7.22218E+01	1.31372E+00	7.08533E+01
BayesExpGiven.RmsAmplitude_2	7.19484E+01	1.22696E+00	7.06506E+01
BayesExpGiven.RmsAmpTotal	1.01944E+02	1.79284E+00	1.00058E+02

Figure 3.29: The Bayes.Condensed.File is shown here. This file is a condensed version of the outputs found in the Mcmc.Values report. The file consists of one output line for each output parameter including derived parameters. Each output line consists of the parameter name, the mean value computed and standard deviation of the samples gathered in the Markov chain Monte Carlo simulation. The peak parameter value is the value of the parameter in the simulation that had peak posterior probability. The heading line shown in this Figure is not present in the condensed file, it is here only to aid in describing the parameters

probability model is a file containing the probability for the model. Its function is similar to that displayed in the Standard output but the the way the file is produced and the format of the file are completely different. In the Bayes Analyze outputs the Probability model file contains one line for each resonance added to the model. These lines contain a description of the model, the logarithm of the posterior probability, the probability gain, and the date and time the model was added. For a complete description of this file, see Subsection 8.5.3.

Bayes.params The bayes.params file is written by the interface and serves as the input parameter file to Bayes Analyze. It contains various parameter settings and the initial model to be processed. The parameter file is divided into three general sections, a header, the global parameters, and the resonance parameters. Each of these three sections is describe in detail in Subsections 8.3, 8.5.1.2 and 8.5.1.3.

Console.log is a running history of what model is being analyzed at the current time. Here is a small snippet of this file:

```
bayes_analyze (V01.20-00)
Developed by Washington University School of Chemistry and
Monsanto St. Louis NMR Center
Base 10 Log Evidence for The First Resonance is: 268.3

Beginning a   1 resonance model
Base 10 Log Of The Probability for 1 Resonance =-3.75974811E+03
Base 10 Log Evidence for The Next Resonance is: 59.2

Beginning a   2 resonance model
Base 10 Log Of The Probability for  2 Resonances =-3.71837136E+03
Base 10 Log Evidence for The Next Resonance is: 65.6
```

As you can see from the above list, the console log is just an indication of the current model, the logarithm of the posterior probability for that model, and finally the log of the evidence that there is another resonance in the data.

log is a complete list of all of the steps taken in the Levenberg-Marquardt searching algorithm. For a complete description of the log file, see Subsection 8.5.4.

Output the output file is a detailed output from every model processed by Bayes Analyze. For a complete description of the output file, see Subsections 8.5.2

Model the Bayes Analyze model file is used as input to Bayes Model. Bayes Model takes the parameters in the model file and generated an fid model of the Marquardt fid. For a complete description of the model file, see Subsections 8.5.5.1

Status While Bayes Analyze is running it updates a status file with some information about what it is currently doing. This information is written into a status file that is fetched by the status button. For a complete list of the various status messages see Subsections 8.6

Summary 1 When the summary 1 report is run, it goes into the Bayes Output file and locates the model which had the highest posterior probability and then writes that model into the summary1 file. For a complete description of this file, see Sections 8.5.6.

Summary 2 When the summary 2 report is run, it goes into the Bayes Output and Bayes model files and locates the model that had maximum posterior probability and then produces a summary of the report. For a complete description of the summary 2 file, see Sections 8.5.7.

Summary 3 or the regions report is produced whenever a regions file is present in the BayesAnalysisFiles directory. When present the scripts will run the summary 3 report. For a description of how to generate a regions file, see Subsections 3.4.2. For a complete description of the summary3 report, see Sections 8.5.8.

Regions will list the regions file if it exists. For a description of how to generate a regions file, see Subsections 3.4.2. For a complete description of the regions file, see Sections 8.5.8.

3.4.9 Files Viewer

The Files Viewer is a tool provided to assist you in finding files. When activated the viewer opens in your current Bayes Home directory and it shows you a listing of all files and directories in Bayes Home. Using a single mouse click on a file will display that file; while using a double mouse click on a directory will expand that directory. You can use this tool to quickly locate and display files. Note that as of this writing a variable length font is in use by this viewer and consequently, it does not preserver line spacing when a file is displayed.

3.4.10 Fortran/C Code Viewer

The Fortran/C Code Viewer is shown in Fig. 3.30, it is used to view and or modify the Ascii models you have loaded. When a model is loaded using either the "System" or "User" buttons a local copy of the model is stored in the BayesAsciiModels subdirectory of your current Bayes home directory and it is displayed in in the Fortran/C Model Viewer along with a list of all of the currently loaded models, left panel. This viewer contains 5 total widgets, and here is a description of them:

Ascii Models is a list of all of the currently loaded Models. This list is a selection menu and by left mouse clicking on a model, the model will be displayed by the Fortran/C Model Viewer. If multiple models are loaded, then clicking on each model displays that model.

Remove Selected Model will delete the model from you Bayes Home/BayesAsciiModels Subdirectory. Note that when system models are loaded, they are copied to the BayesAsciiModels Subdirectory and removing them will remove them from the Subdirectory, it will not remove them from your system directory.

Edit/Create New Model will open the current model in an editor and allow you to make changes to the model. The modified model can be saved using the current name or it can be renamed. For a description of how to write Fortran/C models, see Section E.

Code will display the source code of the currently selected model. Note in this viewer neither the model nor the parameters is editable. However, the priors can be changed on the Prior Viewer

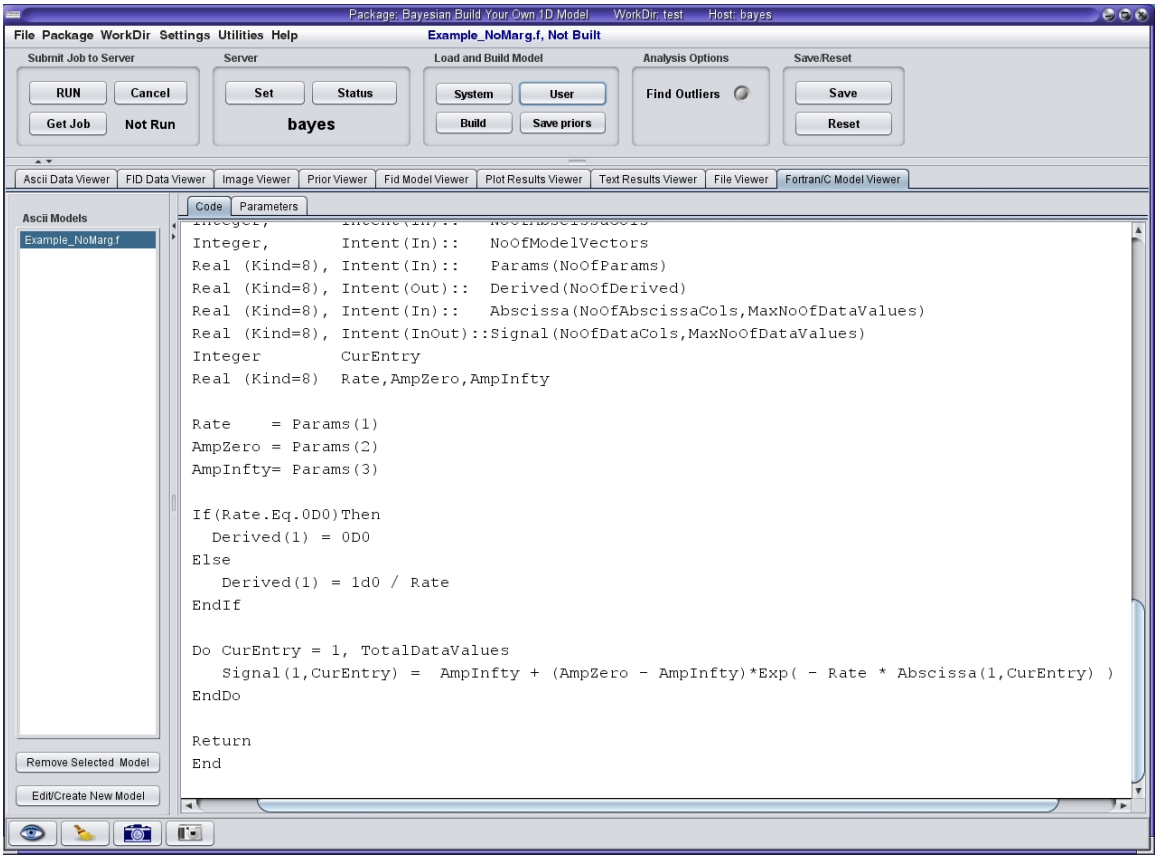

Figure 3.30: The Fortran/C Model Viewer shown here is used to view the models you have currently loaded. You can change a model by selecting it from the "Ascii models" list on the left. The two buttons "Code" and "Parameters" will display the code or the parameters respectively. As a reminder the code is either the Fortran or C code used to implement a model and the parameters file contains essentially a description of the prior probabilities for the model. The display produced by these two buttons are simple displays and neither the code nor the parameters are changeable on this display. However, when the "Edit/Create New Model" button is activated The code is displayed in a popup in a popup editing window. Finally, the "Remove Selected Model" can be used to delete the model from the your BayesAsciiModels subdirectory. Note that activating this button removes your local copy of the model, it does not remove the model from the system directory.

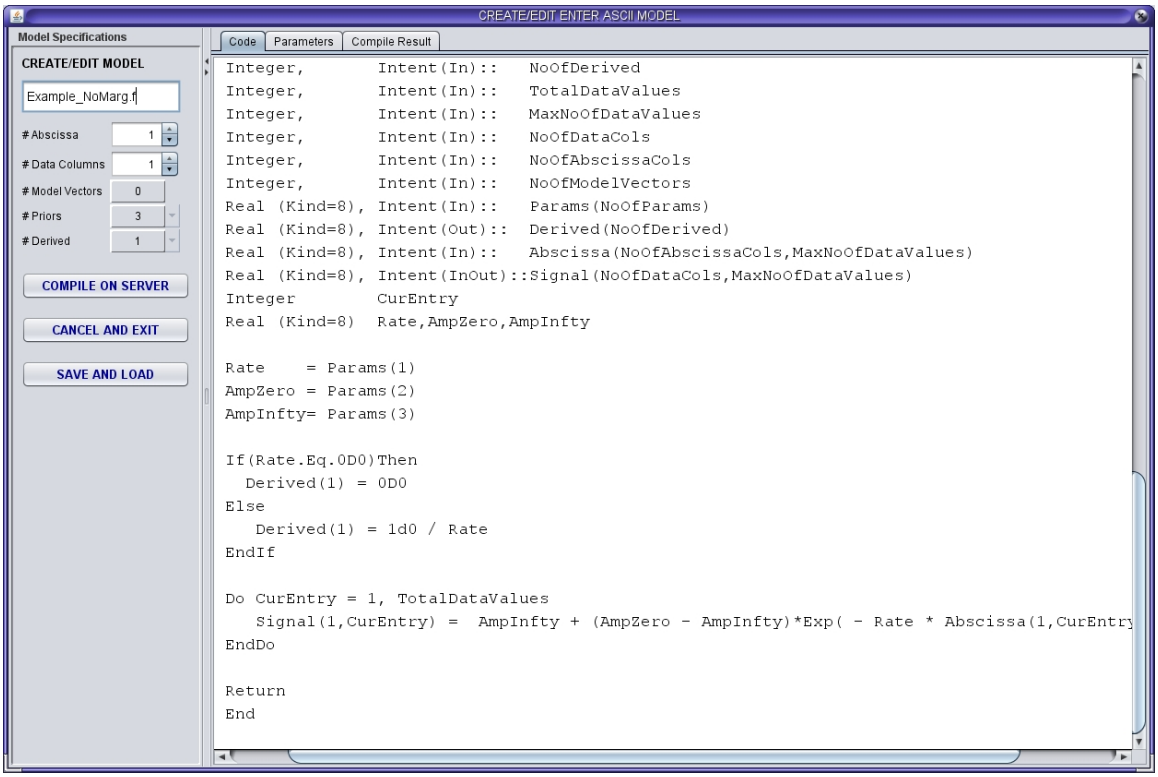

Figure 3.31: The Fortran/C Code Editor shown here is used to edit your currently selected model. In this popup window the text can be modified in any way you please. You can add parameters, change the number of model vectors as well as edit the parameters and the model.

and if you activate the Edit/Create New Models button both the model and the parameters can be modified in the popup window. For a description of how to write Fortran/C models, see Section E.

Parameters will display the parameters file associated with this model. See the above comment on the code.

3.4.10.1 Fortran/C Model Viewer Popup Editor

To edit a model one activates the "Edit/Create New Model" button in the lower left-hand part of the Fortran/C model viewer. When activated the code is copied into a work file and that file is displayed in the Fortran/C model editor shown in Fig 3.4.10.1. On this edit window you can modify the number of parameters, delete/add derived and change the number of model vectors. Here is a rough description of how to do these things:

Create/Edit Model contains the name of the current model. If you wish to create a new model, simply change the name of the model in this field. To save this model you must hit the "Save

and Load" button.

Abscissa can be used to set the number of abscissa columns. Remember if you change the number of abscissa columns the code will probably need to be modified to accommodate this change.

Data Columns can be used to change the number of data columns. Remember if you change the number of data columns the code will need to be modified to accommodate this change.

Model Vectors can only be changed indirectly. This number is a count of the number of parameters having a parameter type of "Amplitude". Consequently, to change the number of model vectors you must change the number of amplitudes.

Priors is a display text that indicates the current number of priors. To change the number of priors activate the grayed out down arrow on the right side of the prior count label. When activated this pull down menu has three options:

Add New Parameter when activated will popup a window asking you for the name of the new parameter. Enter the name of the parameter in the popup and hit OK. The parameter will be added to the list of parameters in the model. However, the prior will be filled in with zero and you must set the prior accordingly. To set the prior simply click on the prior name in the list of priors, this will display the prior in the area on the top of this window. Change the values as you see fit. Note the values are updated as you make the changes.

Remove All will remove all of the priors from this model. Note you are not prompted to prevent you from making a mistake, the priors are simply removed. If this is an error, then hit "Cancel and Exit" to abort your efforts and try again.

Remove parameter will display a selection list and you may pick out the individual prior to be removed.

Derived widget is identical in its behavior as the "Priors" widget. The derived label contains the current number of derived parameters and the grayed out down arrow on the right size allows you to add, remove and remove derived parameters.

Compile On Server will send the current version of the model to the server to compile. Note that this is a simple compile and it does not save the model in your BayesAsciiModels subdirectory.

Cancel And Exit will cancel you current modifications and the Create Enter Ascii model window will exit. Note that if you had activated the "Save and Load" button those changes will still be in effect when cancel is activated.

Save and Load will save your current modifications and the reload the model so that when you exit the popup model editing window your changes are ready to be run.

Along the top of the popup Edit/Create Ascii model are three additional widgets. here is a description of their function

Code will display the current code including any modifications you have made to the code. The window displaying the code is a simple editor and you can make changes to the code as you see fit.

Parameters will display the current parameter file any modifications you have made to the parameters.. The window displaying the code is a simple editor and you can make changes to the code as you see fit. The window displaying the parameters is not quite a edit window, however you can use the parameter window to edit the parameters. First, you can change the parameter being displayed along the top in two ways. You can simply left-mouse click on the parameter you wish to view or you the down arrow on the right-hand side of the displayed parameter name can be used to select a parameter. Regardless, after selecting a parameter, you can use change any value displayed in the top line. As you change these values the parameter file is automatically updated. There are about nine different fields displayed for each parameter, here is a brief list of their functions:

Name is the name of the current parameter. You can simply enter another name if you wish to change this. Note that names must be unique within a model.

Low is the lowest value the parameter can take on. Simply enter any value you wish to change the low value. Note that you can enter any value in this low field, including something larger than the high and the popup will not complain. However, such errors must be corrected before you will be able to save the prior.

Mean is the mean value of the Gaussian prior, and this filed is also used as the peak value in a positive prior. Note that this editor does not change the labels on the prior fields when the priors types are changed.

High is the highest value this parameter can take on and is used by all priors.

Sdev is an abbreviation for Standard Deviation and is used on Gaussian Prior probabilities.

Prior Type is a pull down menu that allows you to select the prior type.

Edit (type) sets whether or not the prior type can be edited by the user. The default allows the field to be edited.

Order indicates if this parameter must be ordered or not. The default is not to order the parameters. If the parameters are to be ordered, they can be ordered from low to high or high to low.

Edit (order) indicates if the order parameter can be modified or not. Usually when a parameter is ordered, this cannot be safely change. Safely change in the sense that the model will continue to work correctly. Consequently, we default this to not editable, but allow the user to change this field at his own peril.

Non-Linear is the default parameter type, this selection menu can be used to set the parameter to amplitude or parameter.

Compile Results will redisplay the results from the previous compile.

Bibliography

[1] Bayes, Rev. T. (1763), "An Essay Toward Solving a Problem in the Doctrine of Chances," *Philos. Trans. R. Soc. London* **53,** pp. 370-418; reprinted in *Biometrika* **45,** pp. 293-315 (1958), and *Facsimiles of Two Papers by Bayes,* with commentary by W. Edwards Deming, New York, Hafner, 1963.

[2] Bretthorst, G. Larry (1988), "Bayesian Spectrum Analysis and Parameter Estimation," in *Lecture Notes in Statistics,* **48**, J. Berger, S. Fienberg, J. Gani, K. Krickenberg, and B. Singer (eds), Springer-Verlag, New York, New York.

[3] Bretthorst, G. Larry (1990), "An Introduction to Parameter Estimation Using Bayesian Probability Theory," in *Maximum Entropy and Bayesian Methods,* Dartmouth College 1989, P. Fougère *ed.,* Kluwer Academic Publishers, Dordrecht the Netherlands, pp. 53-79.

[4] Bretthorst, G. Larry (1990), "Bayesian Analysis I. Parameter Estimation Using Quadrature NMR Models" *J. Magn. Reson.,* **88**, pp. 533-551.

[5] Bretthorst, G. Larry (1990), "Bayesian Analysis II. Signal Detection And Model Selection" *J. Magn. Reson.,* **88**, pp. 552-570.

[6] Bretthorst, G. Larry (1990), "Bayesian Analysis III. Examples Relevant to NMR" *J. Magn. Reson.,* **88**, pp. 571-595.

[7] Bretthorst, G. Larry (1991), "Bayesian Analysis. IV. Noise and Computing Time Considerations," *J. Magn. Reson.,* **93**, pp. 369-394.

[8] Bretthorst, G. Larry (1992), "Bayesian Analysis. V. Amplitude Estimation for Multiple Well-Separated Sinusoids," *J. Magn. Reson.,* **98**, pp. 501-523.

[9] Bretthorst, G. Larry (1992), "Estimating The Ratio Of Two Amplitudes In Nuclear Magnetic Resonance Data," in *Maximum Entropy and Bayesian Methods,* C. R. Smith *et al.* (eds.), pp. 67-77, Kluwer Academic Publishers, the Netherlands.

[10] Bretthorst, G. Larry, (1993), "On The Difference In Means," in Physics & Probability Essays in honor of Edwin T. Jaynes, W. T. Grandy and P. W. Milonni (eds.), pp. 177-194, Cambridge University Press, England.

[11] Bretthorst, G. Larry (1996), "An Introduction To Model Selection Using Bayesian Probability Theory," in *Maximum Entropy and Bayesian Methods*, G. R. Heidbreder, ed., pp. 1-42, Kluwer Academic Publishers, Printed in the Netherlands.

[12] Bretthorst, G. Larry (1999), "The Near-Irrelevance of Sampling Frequency Distributions," in *Maximum Entropy and Bayesian Methods,* W. von der Linden *et al.* (eds.), pp. 21-46, Kluwer Academic Publishers, the Netherlands.

[13] Bretthorst, G. Larry (2001), "Nonuniform Sampling: Bandwidth and Aliasing," in *Maximum Entropy and Bayesian Methods in Science and Engineering,* Joshua Rychert, Gary Erickson and C. Ray Smith *eds.*, pp. 1-28, American Institute of Physics, USA.

[14] Bretthorst, G. Larry, Christopher D. Kroenke, and Jeffrey J. Neil (2004), "Characterizing Water Diffusion In Fixed Baboon Brain," in Bayesian Inference And Maximum Entropy Methods In Science And Engineering, Rainer Fischer, Roland Preuss and Udo von Toussaint *eds.*, AIP conference Proceedings **735**, pp. 3-15.

[15] Bretthorst, G. Larry William C. Hutton, Joel R. Garbow, Joseph J.H. Ackerman, (2005), "Exponential parameter estimation (in NMR) using Bayesian probability theory," *Concepts in Magnetic Resonance,* 27A, Issue 2, pp. 55-63.

[16] Bretthorst, G. Larry William C. Hutton, Joel R. Garbow, Joseph J.H. Ackerman, (2005), "Exponential model selection (in NMR) using Bayesian probability theory," *Concepts in Magnetic Resonance,* 27A, Issue 2, pp. 64-72.

[17] Bretthorst, G. Larry, William C. Hutton, Joel R. Garbow, Joseph J.H. Ackerman, (2005), "How accurately can parameters from exponential models be estimated? A Bayesian view," *Concepts in Magnetic Resonance,* 27A, Issue 2, pp. 73-83.

[18] Bretthorst, G. Larry, W. C. Hutton, J. R. Garbow, J. J. H. Ackerman, (2008), "High Dynamic Range MRS Time-Domain Signal Analysis," *Magn. Reson. in Med.*, **62**, pp. 1026-1035.

[19] Chandramouli, Visvanathan, Karin Ekberg, William C. Schumann, Satish C. Kalhan, John Wahren, and Bernard R. Landau (1997), "Quantifying gluconeogenesis during fasting," *American Journal of Physiology,* **273,** pp. H1209-H1215.

[20] Cox R. T. (1961), "The Algebra of Probable Inference," Johns Hopkins Univ. Press, Baltimore.

[21] d'Avignon, André G. Larry Bretthorst, Marlyn Emerson Holtzer, and Alfred Holtzer (1998), "Site-Specific Thermodynamics and Kinetics of a Coiled-Coil Transiton by Spin Inversion Transfer NMR," *Biophysical Journal,* **74**, pp. 3190-3197.

[22] d'Avignon, André G. Larry Bretthorst, Marlyn Emerson Holtzer, and Alfred Holtzer, (1999), "Thermodynamics and Kinetics of a Folded-Folded Transition at Valine-9 of a GCN4-Like Leucine Zipper," *Biophysical Journal,* **76**, pp. 2752-2759.

[23] Gilks, W. R., S. Richardson and D. J. Spiegelhalter (1996), "Markov Chain Monte Carlo in Practice," Chapman & Hall, London.

[24] Goggans, Paul M. and Ying Chi (2004), "Using Thermodynamic Integration to Calculate the Posterior Probability in Bayesian Model Selection Problems," in *Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 23rd International Workshop,* Volume 707, pp. 59-66.

[25] Holtzer, Marlyn Emerson, G. Larry Bretthorst, D. André d'Avignon, Ruth Hogue Angelette, Lisa Mints, and Alfred Holtzer (2001), "Temperature Dependence of the Folding and Unfolding Kinetics of the GCN4 Leucine Lipper via 13C alpha-NMR,'' *Biophysical Journal,* **80,** pp. 939-951.

[26] Jaynes, E. T. (1968), "Prior Probabilities," *IEEE Transactions on Systems Science and Cybernetics,* SSC-4, pp. 227-241; reprinted in [29].

[27] Jaynes, E. T. (1978), "Where Do We Stand On Maximum Entropy?" in *The Maximum Entropy Formalism,* R. D. Levine and M. Tribus *Eds.,* pp. 15-118, Cambridge: MIT Press, Reprinted in [29].

[28] Jaynes, E. T. (1980), "Marginalization and Prior Probabilities," in *Bayesian Analysis in Econometrics and Statistics,* A. Zellner, *ed.,* North-Holland Publishing Company, Amsterdam; reprinted in [29].

[29] Jaynes, E. T. (1983), "Papers on Probability, Statistics and Statistical Physics," a reprint collection, D. Reidel, Dordrecht the Netherlands; second edition Kluwer Academic Publishers, Dordrecht the Netherlands, 1989.

[30] Jaynes, E. T. (1957), "How Does the Brain do Plausible Reasoning?" unpublished Stanford University Microwave Laboratory Report No. 421; reprinted in *Maximum-Entropy and Bayesian Methods in Science and Engineering* **1,** pp. 1-24, G. J. Erickson and C. R. Smith *Eds.,* 1988.

[31] Jaynes, E. T. (2003), "Probability Theory—The Logic of Science," edited by G. Larry Bretthorst, Cambridge University Press, Cambridge UK.

[32] Jeffreys, Harold Sir (1939), "Theory of Probability," Oxford Univ. Press, London; Later editions, 1948, 1961.

[33] Jones, John G. (2001), Michael A. Solomon, Suzanne M. Cole, A. Dean Sherry, Craig R. Malloy "An integrated ^{2}H and ^{13}C NMR study of gluconeogenesis and TCA cycle flux in humans," *American Journal of Physiology, Endocrinology, and Metabolism,* **281,** pp. H848-H856.

[34] Kotyk, John, N. G. Hoffman, W. C. Hutton, G. Larry Bretthorst, and J. J. H. Ackerman (1992), "Comparison of Fourier and Bayesian Analysis of NMR Signals. I. Well-Separated Resonances (The Single-Frequency Case)," *J. Magn. Reson.*, **98**, pp. 483–500.

[35] Laplace, Pierre Simon (1814), "A Philosophical Essay on Probabilities," John Wiley & Sons, London, Chapman & Hall, Limited 1902. Translated from the 6th edition by F. W. Truscott and F. L. Emory.

[36] Lartillot, N. and H. Philippe (2006), "Computing Bayes Factors Using Thermodynamic Integration," *Systematic Biology*, 55(2) pp. 195-207.

[37] Le Bihan, D. (1985), E. Breton, "Imagerie de diffusion in-vivo par rsonance," C R Acad Sci (Paris) **301** (15) pp. 1109-1112.

[38] Lomb, N. R. (1976), "Least-Squares Frequency Analysis of Unevenly Spaced Data," *Astrophysical and Space Science*, **39**, pp. 447-462.

[39] Loredo, T. J. (1990), "From Laplace To SN 1987A: Bayesian Inference In Astrophysics," in *Maximum Entropy and Bayesian Methods,* P. F. Fougere (ed), Kluwer Academic Publishers, Dordrecht, The Netherlands.

[40] Malloy, Craig R., A. Dean Sherry, F. Mark H. Jeffrey (1988), "Evaluation of Carbon Flux and Substrate Selection through Alternate Pathways Involving the Citric Acid Cycle of the Heart by 13C NMR Spectroscopy," *Journal of Biological Chemistry,* Vol. 263, No. 15, pp. 6964-6971.

[41] Malloy, Craig R., A. Dean Sherry, F. Mark H. Jeffrey (1990), "Analysis of tricarboxylic acid cycle of the heart using ^{13}C isotope isomers," *American Journal of Physiology,* **259,** pp. H987-H995.

[42] Merboldt, K., Hanicke, W., Frahm, J. (1969), "Self-diffusion NMR imaging using stimulated echoes," Journal of Magnetic Resonance **64** (3) pp. 479-486.

[43] Metropolis, Nicholas, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller (1953), "Equation of State Calculations by Fast Computing Machines," Journal of Chemical Physics. The previous link is to the Americain Institute of Physics and if you do not have access to Science Sitations you many not be able to retrieve this paper.

[44] Neal, Radford M. (1993), "Probabilistic Inference Using Markov Chain Monte Carlo Methods," technical report CRG-TR-93-1, Dept. of Computer Science, University of Toronto.

[45] Neil, Jeffrey J., and G. Larry Bretthorst (1993), "On the Use of Bayesian Probability Theory for Analysis of Exponential Decay Data: An Example Taken from Intravoxel Incoherent Motion Experiments," *Magn. Reson. in Med.*, **29**, pp. 642–647.

[46] Nyquist, H. (1924), "Certain Factors Affecting Telegraph Speed," *Bell System Technical Journal,* **3**, pp. 324-346.

[47] Nyquist, H., (1928), "Certain Topics in Telegraph Transmission Theory," *Transactions AIEE*, **3**, p. 617-644.

[48] Press W. H., S. A. Teukolsky, W. T. Vetterling and B. P. Flannary (1992), "Numerical Recipes The Art of Scientific Computing Second Edition," Cambridge University Press, Cambridge UK.

[49] Scargle, J. D. (1982), "Studies in Astronomical Time Series Analysis II. Statistical Aspects of Spectral Analysis of Unevenly Sampled Data," *Astrophysical Journal*, **263**, pp. 835-853.

[50] Scargle, J. D. (1989), "Studies in Astronomical Time Series Analysis. III. Fourier Transforms, Autocorrelation and Cross-correlation Functions of Unevenly Spaced Data," *Astrophysical Journal*, **343**, pp. 874-887.

[51] Schuster, A., (1905), "The Periodogram and its Optical Analogy," *Proceedings of the Royal Society of London,* **77**, p. 136-140.

[52] Shannon, C. E. (1948), "A Mathematical Theory of Communication," *Bell Syst. Tech. J.* **27,** pp. 379-423.

[53] Shore J. E., R. W. Johnson (1981), "Properties of cross-entropy minimization," *IEEE Trans. on Information Theory,* **IT-27,** No. 4, pp. 472-482.

[54] Shore J. E., R. W. Johnson (1980), "Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy," *IEEE Trans. on Information Theory,* **IT-26,** No. 1, pp. 26-37.

[55] Sivia, D. S. and J. Skilling (2006), "Data Analysis: A Bayesian Tutorial," Oxford University Press, USA.

[56] Stejskal, E. O., Tanner, J. E. (1965), "Spin Diffusion Measurements: Spin Echoes in the Presence of a Time-Dependent Field Gradient." Journal of Chemical Physics **42** (1), pp. 288-292.

[57] Taylor, D. G., Bushell, M. C. (1985), "The spatial mapping of translational diffusion coefficients by the NMR imaging technique," Physics in Medicine and Biology **30** (4), pp. 345-349.

[58] Tribus, M. (1969), "Rational Descriptions, Decisions and Designs," Pergamon Press, Oxford.

[59] Woodward, P. M. (1953), "Probability and Information Theory, with Applications to Radar," McGraw-Hill, N. Y. Second edition (1987); R. E. Krieger Pub. Co., Malabar, Florida. 1990.

[60] Zellner, A. (1971), "An Introduction to Bayesian Inference in Econometrics," John Wiley and Sons, New York.